

DEEP NETWORKS FROM THE PRINCIPLE OF RATE REDUCTION

Kwan Ho Ryan Chan[†] Yaodong Yu[†] Chong You[†] Haozhi Qi[†] John Wright^{‡◊} Yi Ma[†][†]Department of EECS, University of California, Berkeley[‡]Department of Electrical Engineering and Data Science Institute, Columbia University[◊]Department of Applied Physics and Applied Mathematics, Columbia University

ABSTRACT

This work attempts to interpret modern deep (convolutional) networks from the principles of rate reduction and (shift) invariant classification. We show that the basic iterative gradient ascent scheme for optimizing the rate reduction of learned features naturally leads to a multi-layer deep network, one iteration per layer. The layered architectures, linear and nonlinear operators, and even parameters of the network are all explicitly constructed layer-by-layer in a forward propagation fashion by emulating the gradient scheme. All components of this “white box” network have precise optimization, statistical, and geometric interpretation. This principled framework also reveals and justifies the role of multi-channel lifting and sparse coding in early stage of deep networks. Moreover, all linear operators of the so-derived network naturally become multi-channel convolutions when we enforce classification to be rigorously shift-invariant. The derivation also indicates that such a convolutional network is significantly more efficient to construct and learn in the spectral domain. Our preliminary simulations and experiments indicate that so constructed deep network can already learn a good discriminative representation even without any back propagation training.

1 INTRODUCTION AND MOTIVATION

In recent years, various deep (convolution) network architectures such as AlexNet (Krizhevsky et al., 2012), VGG (Simonyan & Zisserman, 2015), ResNet (He et al., 2016), DenseNet (Huang et al., 2017), Recurrent CNN, LSTM (Hochreiter & Schmidhuber, 1997), Capsule Networks (Hinton et al., 2011), etc., have demonstrated very good performance in classification tasks of real-world datasets such as speeches or images. Nevertheless, almost all such networks are developed through years of empirical *trial and error*, including both their architectures/operators and the ways they are to be effectively trained. Some recent practices even take to the extreme by searching for effective network structures and training strategies through extensive random search techniques, such as Neural Architecture Search (Zoph & Le, 2017; Baker et al., 2017), AutoML (Hutter et al., 2019), and Learning to Learn (Andrychowicz et al., 2016).

Despite tremendous empirical advances, there is still a lack of rigorous theoretical justification of the need or reasons for “deep” network architectures and a lack of fundamental understanding of the associated operators (e.g. multi-channel convolution and nonlinear activation) in each layer. As a result, deep networks are often designed and trained heuristically and then used as a “black box.” There have been a severe lack of guiding principles for each of the stages: For a given task, how wide or deep the network should be? What are the roles and relationships among the multiple (convolution) channels? Which parts of the networks need to be learned and trained and which can be determined in advance? How to evaluate the optimality of the resulting network? As a consequence, besides empirical evaluation, it is usually impossible to offer any rigorous guarantees for certain performance of a trained network, such as invariance to transformation (Azulay & Weiss, 2018; Engstrom et al., 2017) or overfitting noisy or even arbitrary labels (Zhang et al., 2017).

*The first three authors contributed equally to this work.

In this paper, we do not intend to address all these questions but we would attempt to offer a plausible interpretation of deep (convolution) neural networks by deriving a class of deep networks *from first principles*. We contend that all key features and structures of modern deep (convolution) neural networks can be naturally derived from optimizing a principled objective, namely the *rate reduction* recently proposed by Yu et al. (2020), that seeks a compact discriminative (invariant) representation of the data. More specifically, the basic iterative *gradient ascent* scheme for optimizing the objective naturally takes the form of a deep neural network, one layer per iteration.

This principled approach brings a couple of nice surprises: First, architectures, operators, and parameters of the network can be constructed explicitly layer-by-layer in a *forward propagation* fashion, and all inherit precise optimization, statistical and geometric interpretation. As result, the so constructed “white box” deep network already gives a good discriminative representation (and achieves good classification performance) *without any back propagation* for training the deep network. Second, in the case of seeking a representation *rigorously* invariant to shift or translation, the network naturally lends itself to a multi-channel convolutional network. Moreover, the derivation indicates such a convolutional network is computationally more efficient to learn and construct in the *spectral (Fourier) domain*, analogous to how neurons in the visual cortex encode and transit information with their spiking frequencies (Eliasmith & Anderson, 2003; Belitski et al., 2008).

2 TECHNICAL APPROACH

Consider a basic classification task: given a set of m samples $\mathbf{X} \doteq [\mathbf{x}^1, \dots, \mathbf{x}^m] \in \mathbb{R}^{n \times m}$ and their associated memberships $\pi(\mathbf{x}^i) \in [k]$ in k different classes, a deep network is typically used to model a direct mapping from the input data $\mathbf{x} \in \mathbb{R}^n$ to its class label $f(\mathbf{x}, \boldsymbol{\theta}) : \mathbf{x} \mapsto \mathbf{y} \in \mathbb{R}^k$, where \mathbf{y} is typically a “one-hot” vector encoding the membership information $\pi(\mathbf{x})$: the j -th entry of \mathbf{y} is 1 iff $\pi(\mathbf{x}) = j$. The parameters $\boldsymbol{\theta}$ of the network is typically learned to minimize certain prediction loss, say the cross entropy loss, via gradient-descent type back propagation. Although this popular approach provides people a direct and effective way to train a network that predicts the class information, the so learned representation is however implicit and lacks clear interpretation.

2.1 PRINCIPLE OF RATE REDUCTION AND GROUP INVARIANCE

The Principle of Maximal Coding Rate Reduction. To help better understand features learned in a deep network, the recent work of Yu et al. (2020) has argued that the goal of (deep) learning is to learn a compact discriminative and diverse feature representation¹ $\mathbf{z} = f(\mathbf{x}) \in \mathbb{R}^n$ of the data \mathbf{x} before any subsequent tasks such as classification: $\mathbf{x} \xrightarrow{f(\mathbf{x})} \mathbf{z} \xrightarrow{h(\mathbf{z})} \mathbf{y}$. To be more precise, instead of directly fitting the class label \mathbf{y} , a principled objective is to learn a feature map $f(\mathbf{x}) : \mathbf{x} \mapsto \mathbf{z}$ which transforms the data \mathbf{x} onto a set of most discriminative low-dimensional linear subspaces $\{\mathcal{S}^j\}_{j=1}^k \subset \mathbb{R}^n$, one subspace \mathcal{S}^j per class $j \in [k]$.

Let $\mathbf{Z} \doteq [\mathbf{z}^1, \dots, \mathbf{z}^m] = [f(\mathbf{x}^1), \dots, f(\mathbf{x}^m)]$ be the features of the given samples \mathbf{X} . WLOG, we may assume all features \mathbf{z}^i are normalized to be of unit norm: $\mathbf{z}^i \in \mathbb{S}^{n-1}$. For convenience, let $\mathbf{\Pi}^j \in \mathbb{R}^{m \times m}$ be a diagonal matrix whose diagonal entries encode the membership of samples/features belong to the j -th class: $\mathbf{\Pi}^j(i, i) = \pi(\mathbf{x}^i) = \pi(\mathbf{z}^i)$. Then based on principles from lossy data compression (Ma et al., 2007), Yu et al. (2020) has suggested that the optimal representation $\mathbf{Z}_* \subset \mathbb{S}^{n-1}$ should maximize the following coding rate reduction objective, known as the MCR² principle:

$$\text{Rate Reduction: } \Delta R(\mathbf{Z}) \doteq \underbrace{\frac{1}{2} \log \det (\mathbf{I} + \alpha \mathbf{Z} \mathbf{Z}^*)}_{R(\mathbf{Z})} - \underbrace{\sum_{j=1}^k \frac{\gamma_j}{2} \log \det (\mathbf{I} + \alpha_j \mathbf{Z} \mathbf{\Pi}^j \mathbf{Z}^*)}_{R_c(\mathbf{Z}, \mathbf{\Pi})}, \quad (1)$$

where $\alpha = n/(m\epsilon^2)$, $\alpha_j = n/(\text{tr}(\mathbf{\Pi}^j)\epsilon^2)$, $\gamma_j = \text{tr}(\mathbf{\Pi}^j)/m$ for $j = 1, \dots, k$. Given a prescribed quantization error ϵ , the first term R of $\Delta R(\mathbf{Z})$ measures the total coding length for all the features \mathbf{Z} and the second term R_c is the sum of coding lengths for features in each of the k classes.

¹To simplify the presentation, we assume for now that the feature \mathbf{z} and \mathbf{x} have the same dimension n . But in general they can be different as we will soon see, say in the case \mathbf{z} is multi-channel extracted from \mathbf{x} .

In Yu et al. (2020), the authors have shown the optimal representation \mathbf{Z}_* that maximizes the above objective indeed has desirable nice properties. Nevertheless, they adopted a conventional deep network (e.g. ResNet) as a black box to model and parameterize the feature mapping: $\mathbf{z} = f(\mathbf{x}, \boldsymbol{\theta})$. It has empirically shown that with such a choice, one can effectively optimize the MCR² objective and obtain discriminative and diverse representations for classifying real image data.

However, there remain several unanswered problems. Although the resulting feature representation is more interpretable, the network itself is still not. It is *not* clear why any chosen network is able to optimize the desired MCR² objective: Would there be any potential limitations? The good empirical results (say with a ResNet) do not necessarily justify the particular choice in architectures and operators of the network: Why is a layered model necessary, how wide and deep is adequate, and is there any rigorous justification for the use of convolutions and nonlinear operators used? In Section 2.2, we show that using gradient ascent to maximize the rate reduction $\Delta R(\mathbf{Z})$ naturally leads to a “white box” deep network that represents such a mapping. All linear/nonlinear operators and parameters of the network are *explicitly constructed in a purely forward propagation fashion*.

Group Invariant Rate Reduction. So far, we have considered the data and features as vectors. In many applications, such as serial data or imagery data, the semantic meaning (labels) of the data and their features are *invariant* to certain transformations $\mathbf{g} \in \mathbb{G}$ (for some group \mathbb{G}) (Cohen & Welling, 2016). For example, the meaning of an audio signal is invariant to shift in time; and the identity of an object in an image is invariant to translation in the image plane. Hence, we prefer the feature mapping $f(\mathbf{x}, \boldsymbol{\theta})$ is rigorously invariant to such transformations:

$$\text{Group Invariance: } f(\mathbf{x} \circ \mathbf{g}, \boldsymbol{\theta}) \sim f(\mathbf{x}, \boldsymbol{\theta}), \quad \forall \mathbf{g} \in \mathbb{G}, \quad (2)$$

where “ \sim ” indicates two features belonging to the same equivalent class. The recent work of Zaheer et al. (2017); Maron et al. (2020) characterize properties of networks and operators for set permutation groups. Nevertheless, it remains challenging to learn features via a deep network that are *guaranteed* to be invariant even to simple transformations such as translation and rotation (Azulay & Weiss, 2018; Engstrom et al., 2017). In Section 2.3, we show that the MCR² principle is compatible with invariance in a very natural and precise way: we only need to assign all transformed versions $\{\mathbf{x} \circ \mathbf{g} \mid \mathbf{g} \in \mathbb{G}\}$ into the same class as \mathbf{x} and map them all to the same subspace \mathcal{S} .² We will rigorously show (in the Appendices) that, when the group \mathbb{G} is (discrete) circular 1D shifting or 2D translation, the resulting deep network naturally becomes a *multi-channel convolution network*!

2.2 DEEP NETWORKS FROM OPTIMIZING RATE REDUCTION

Gradient Ascent for Rate Reduction on the Training Samples. First let us directly try to optimize the objective $\Delta R(\mathbf{Z})$ as a function in the training samples $\mathbf{Z} \subset \mathbb{S}^{n-1}$. To this end, we may adopt a (projected) *gradient ascent* scheme, for some step size $\eta > 0$:

$$\mathbf{Z}_{\ell+1} \propto \mathbf{Z}_{\ell} + \eta \cdot \left. \frac{\partial \Delta R}{\partial \mathbf{Z}} \right|_{\mathbf{Z}_{\ell}} \quad \text{subject to } \mathbf{Z}_{\ell+1} \subset \mathbb{S}^{n-1}. \quad (3)$$

This scheme can be interpreted as how one should incrementally adjust locations of the current features \mathbf{Z}_{ℓ} in order for the resulting $\mathbf{Z}_{\ell+1}$ to improve the rate reduction $\Delta R(\mathbf{Z})$. Simple calculation shows that the gradient $\frac{\partial \Delta R}{\partial \mathbf{Z}}$ entails evaluating the following derivatives of the terms in (1):

$$\left. \frac{1}{2} \frac{\partial \log \det(\mathbf{I} + \alpha \mathbf{Z} \mathbf{Z}^*)}{\partial \mathbf{Z}} \right|_{\mathbf{Z}_{\ell}} = \underbrace{\alpha (\mathbf{I} + \alpha \mathbf{Z}_{\ell} \mathbf{Z}_{\ell}^*)^{-1} \mathbf{Z}_{\ell}}_{\mathbf{E}_{\ell} \in \mathbb{R}^{n \times n}} \in \mathbb{R}^{n \times m}, \quad (4)$$

$$\left. \frac{1}{2} \frac{\partial (\gamma_j \log \det(\mathbf{I} + \alpha_j \mathbf{Z} \boldsymbol{\Pi}^j \mathbf{Z}^*))}{\partial \mathbf{Z}} \right|_{\mathbf{Z}_{\ell}} = \gamma_j \underbrace{\alpha_j (\mathbf{I} + \alpha_j \mathbf{Z}_{\ell} \boldsymbol{\Pi}^j \mathbf{Z}_{\ell}^*)^{-1} \mathbf{Z}_{\ell} \boldsymbol{\Pi}^j}_{\mathbf{C}_{\ell}^j \in \mathbb{R}^{n \times n}} \in \mathbb{R}^{n \times m}. \quad (5)$$

Notice that in the above, the matrix \mathbf{E}_{ℓ} only depends on \mathbf{Z}_{ℓ} and it aims to *expand* all the features to increase the overall coding rate; the matrix \mathbf{C}_{ℓ}^j depends on features from each class and aims to *compress* them to reduce the coding rate of each class. We provide the geometric and statistic

²Hence, any subsequent classifiers defined on the resulting set of subspaces will be automatically invariant to such transformations.




Figure 1: Interpretation of E_ℓ and C_ℓ^j : E_ℓ expands all features by contrasting and repelling features across different classes; C_ℓ^j compresses each class by contracting the features to a low-dimensional subspace.

meaning of E_ℓ and C_ℓ^j in Remark 1 below. Then the complete gradient $\frac{\partial \Delta R}{\partial \mathbf{Z}} \Big|_{\mathbf{Z}_\ell}$ is of the following form:

$$\frac{\partial \Delta R}{\partial \mathbf{Z}} \Big|_{\mathbf{Z}_\ell} = \underbrace{\mathbf{E}_\ell}_{\text{Expansion}} \mathbf{Z}_\ell - \sum_{j=1}^k \gamma_j \underbrace{\mathbf{C}_\ell^j}_{\text{Compression}} \mathbf{Z}_\ell \mathbf{\Pi}^j \in \mathbb{R}^{n \times m}. \quad (6)$$

Remark 1 (Interpretation of E_ℓ and C_ℓ^j as Linear Operators) For any $\mathbf{z}_\ell \in \mathbb{R}^n$, we have

$$\mathbf{E}_\ell \mathbf{z}_\ell = \alpha(\mathbf{z}_\ell - \mathbf{Z}_\ell \mathbf{q}_\ell^*) \quad \text{where} \quad \mathbf{q}_\ell^* \doteq \underset{\mathbf{q}_\ell}{\operatorname{argmin}} \alpha \|\mathbf{z}_\ell - \mathbf{Z}_\ell \mathbf{q}_\ell\|_2^2 + \|\mathbf{q}_\ell\|_2^2. \quad (7)$$

Notice that \mathbf{q}_ℓ^* is exactly the solution to the ridge regression by all the data points \mathbf{Z}_ℓ concerned. Therefore, E_ℓ (similarly for C_ℓ^j) is approximately (i.e. when m is large enough) the projection onto the orthogonal complement of the subspace spanned by columns of \mathbf{Z}_ℓ . Another way to interpret the matrix E_ℓ is through eigenvalue decomposition of the covariance matrix $\mathbf{Z}_\ell \mathbf{Z}_\ell^*$. Assuming that $\mathbf{Z}_\ell \mathbf{Z}_\ell^* \doteq \mathbf{U}_\ell \mathbf{\Lambda}_\ell \mathbf{U}_\ell^*$ where $\mathbf{\Lambda}_\ell \doteq \operatorname{diag}\{\sigma_1, \dots, \sigma_d\}$, we have

$$\mathbf{E}_\ell = \alpha \mathbf{U}_\ell \operatorname{diag} \left\{ \frac{1}{1 + \alpha \sigma_1}, \dots, \frac{1}{1 + \alpha \sigma_d} \right\} \mathbf{U}_\ell^*. \quad (8)$$

Therefore, the matrix E_ℓ operates on a vector \mathbf{z}_ℓ by stretching in a way that directions of large variance are shrunk while directions of vanishing variance are kept. These are exactly the directions (4) in which we move the features so that the overall volume expands and the coding rate will increase, hence the positive sign. To the opposite effect, the directions associated with (5) are “residuals” of features of each class deviate from the subspace to which they are supposed to belong. These are exactly the directions in which the features need to be compressed back onto their respective subspace, hence the negative sign (see Figure 1).

Essentially, the linear operations E_ℓ and C_ℓ^j in gradient ascent for rate reduction are determined by training data conducting “auto-regressions”. The recent renewed understanding about ridge regression in an over-parameterized setting (Yang et al., 2020; Wu & Xu, 2020) indicates that using seemingly redundantly sampled data (from each subspaces) as regressors do not lead to overfitting.

Gradient-Guided Feature Map Increment. Notice that in the above, the gradient ascent considers all the features $\mathbf{Z}_\ell = [\mathbf{z}_\ell^1, \dots, \mathbf{z}_\ell^m]$ as free variables. The increment $\mathbf{Z}_{\ell+1} - \mathbf{Z}_\ell = \eta \frac{\partial \Delta R}{\partial \mathbf{Z}} \Big|_{\mathbf{Z}_\ell}$ does not yet give a transform on the entire feature domain $\mathbf{z}_\ell \in \mathbb{R}^n$. Hence, in order to find the optimal $f(\mathbf{x}, \boldsymbol{\theta})$ explicitly, we may consider constructing a small increment transform $g(\cdot, \boldsymbol{\theta}_\ell)$ on the ℓ -th layer feature \mathbf{z}_ℓ to emulate the above (projected) gradient scheme:

$$\mathbf{z}_{\ell+1} \propto \mathbf{z}_\ell + \eta \cdot g(\mathbf{z}_\ell, \boldsymbol{\theta}_\ell) \quad \text{subject to} \quad \mathbf{z}_{\ell+1} \in \mathbb{S}^{n-1} \quad (9)$$

such that: $[g(\mathbf{z}_\ell^1, \boldsymbol{\theta}_\ell), \dots, g(\mathbf{z}_\ell^m, \boldsymbol{\theta}_\ell)] \approx \frac{\partial \Delta R}{\partial \mathbf{Z}} \Big|_{\mathbf{Z}_\ell}$. That is, we need to approximate the gradient flow $\frac{\partial \Delta R}{\partial \mathbf{Z}}$ that locally deforms each (training) feature $\{\mathbf{z}_\ell^i\}_{i=1}^m$ with a continuous mapping $g(\mathbf{z})$ defined on the entire feature space $\mathbf{z}_\ell \in \mathbb{R}^n$.

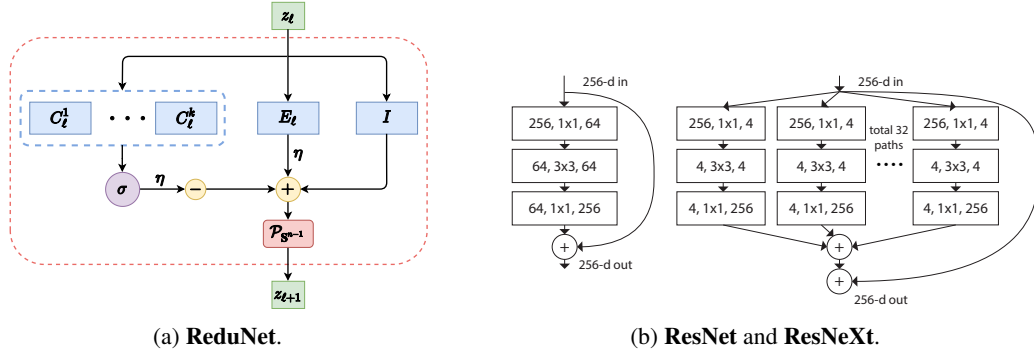


Figure 2: Comparison of Network Architectures. (a): Layer structure of the **ReduNet** derived from one iteration of gradient ascent for optimizing rate reduction. (b) (left): A layer of ResNet (He et al., 2016); and (b) (right): A layer of ResNeXt (Xie et al., 2017). As we will see in the next section, the linear operators E_ℓ and C_ℓ^j of the ReduNet naturally become (multi-channel) convolutions when shift-invariance is imposed.

Remark 2 (Connection and Difference from Neural ODE) One may interpret the increment (9) as a discretized version of a continuous ordinary differential equation (ODE):

$$\dot{z} = g(z, \theta). \quad (10)$$

Hence the (deep) network so constructed can be interpreted as certain neural ODE (Chen et al., 2018). Nevertheless, unlike neural ODE where the flow g is chosen to be some generic structures whose parameters are trained later, here our $g(z, \theta)$ is to emulate the gradient flow of the rate reduction on the feature set

$$\dot{Z} = \eta \cdot \frac{\partial \Delta R}{\partial Z} \implies Z_{\ell+1} \propto Z_\ell + \eta \cdot \frac{\partial \Delta R}{\partial Z} \Big|_{Z_\ell},$$

and its structure and parameters are entirely derived and fully determined from this objective, without any other priors, heuristics or post training.

By inspecting the structure of the gradient (6), it suggests that a natural candidate for the increment transform $g(z_\ell, \theta_\ell)$ is of the form:

$$g(z_\ell, \theta_\ell) \doteq E_\ell z_\ell - \sum_{j=1}^k \gamma_j C_\ell^j z_\ell \pi^j(z_\ell) \in \mathbb{R}^n, \quad (11)$$

where $\pi^j(z_\ell) \in [0, 1]$ indicates the probability of z_ℓ belonging to the j -th class.³ Notice that the increment depends on 1). A linear map represented by E_ℓ that depends only on statistics of all features from the preceding layer; 2). A set of linear maps $\{C_\ell^j\}_{j=1}^k$ and memberships $\{\pi^j(z_\ell)\}_{j=1}^k$ of the features.

Since we only have the membership π^j for the training samples, the function g defined in (11) can only be evaluated on the training samples. To extrapolate the function g to the entire feature space, we need to estimate $\pi^j(z_\ell)$ in its second term. In the conventional deep learning, this map is typically modeled as a deep network and learned from the training data, say via *back propagation*. Nevertheless, our goal here is not to learn a precise classifier $\pi^j(z_\ell)$ already. Instead, we only need a good enough estimate of the class information in order for g to approximate the gradient $\frac{\partial \Delta R}{\partial Z}$ well.

From the geometric interpretation of the linear maps E_ℓ and C_ℓ^j given by Remark 1, the term $p_\ell^j \doteq C_\ell^j z_\ell$ can be viewed as projection of z_ℓ onto the orthogonal complement of each class j . Therefore, $\|p_\ell^j\|_2$ is small if z_ℓ is in class j and large otherwise. This motivates us to estimate its membership based on the following softmax function: $\hat{\pi}^j(z_\ell) \doteq \frac{\exp(-\lambda \|C_\ell^j z_\ell\|)}{\sum_{j=1}^k \exp(-\lambda \|C_\ell^j z_\ell\|)} \in [0, 1]$.

³Notice that on the training samples Z_ℓ , for which the memberships Π^j are known, the so defined $g(z_\ell, \theta)$ gives exactly the values for the gradient $\frac{\partial \Delta R}{\partial Z} \Big|_{Z_\ell}$.

Hence the second term of (11) can be approximated by this estimated membership:⁴

$$\sum_{j=1}^k \gamma_j C_\ell^j z_\ell \pi^j(z_\ell) \approx \sum_{j=1}^k \gamma_j C_\ell^j z_\ell \cdot \hat{\pi}^j(z_\ell) \doteq \sigma([C_\ell^1 z_\ell, \dots, C_\ell^k z_\ell]) \in \mathbb{R}^n, \quad (12)$$

which is denoted as a nonlinear operator $\sigma(\cdot)$ on outputs of the feature z_ℓ through k banks of filters: $[C_\ell^1, \dots, C_\ell^k]$. Notice that the nonlinearity arises due to a “soft” assignment of class membership based on the feature responses from those filters. Overall, combining (9), (11), and (12), the increment feature transform from z_ℓ to $z_{\ell+1}$ now becomes:

$$z_{\ell+1} \propto z_\ell + \eta \cdot E_\ell z_\ell - \eta \cdot \sigma([C_\ell^1 z_\ell, \dots, C_\ell^k z_\ell]) \quad \text{subject to} \quad z_{\ell+1} \in \mathbb{S}^{n-1}, \quad (13)$$

with the nonlinear function $\sigma(\cdot)$ defined above and θ_ℓ collecting all the layer-wise parameters including $E_\ell, C_\ell^j, \gamma_j$ and λ , and with features at each layer always “normalized” onto a sphere \mathbb{S}^{n-1} , denoted as $\mathcal{P}_{\mathbb{S}^{n-1}}$. The form of increment in (13) can be illustrated by a diagram in Figure 2.

Deep Network from Rate Reduction. Notice that the increment is constructed to emulate the gradient ascent for the rate reduction ΔR . Hence by transforming the features iteratively via the above process, we expect the rate reduction to increase, as we will see in the experimental section. This iterative process, once converged say after L iterations, gives the desired feature map $f(x, \theta)$ on the input $z_0 = x$, precisely in the form of a *deep network*, in which each layer has the structure shown in Figure 2:

$$f(x, \theta) = \phi^L \circ \phi^{L-1} \circ \dots \circ \phi^0(x), \quad \text{with} \quad \phi^\ell(z_\ell, \theta_\ell) \doteq \mathcal{P}_{\mathbb{S}^{n-1}}[z_\ell + \eta \cdot g(z_\ell, \theta_\ell)]. \quad (14)$$

As this deep network is derived from maximizing the rate **reduced**, we call it the **ReduNet**. Notice that all parameters of the network are explicitly constructed layer by layer in a *forward propagation* fashion. Once constructed, there is no need of any additional supervised learning, say via back propagation. As suggested in Yu et al. (2020), the so learned features can be directly used for classification via a nearest subspace classifier.

Comparison with Other Approaches and Architectures. Structural similarities between deep networks and iterative optimization schemes, especially those for solving sparse coding, have been long noticed. In particular, Gregor & LeCun (2010) has argued that algorithms for sparse coding, such as the FISTA algorithm (Beck & Teboulle, 2009), can be viewed as a deep network and be trained for better coding performance, known as LISTA. Later Monga et al. (2019); Sun et al. (2020) have proposed similar interpretation of deep networks as unrolling algorithms for sparse coding. Like all networks that are inspired by unfolding certain iterative optimization schemes, the structure of the ReduNet naturally contains a skip connection between adjacent layers as in the ResNet (He et al., 2016). Remark 4 in Appendix A discusses possible improvement to the basic gradient scheme that may introduce additional skip connections beyond adjacent layers. As illustrated in Figure 2, the remaining $k + 1$ parallel channels E, C^j of the ReduNet actually draw resemblance to the parallel structures that people later found empirically beneficial for deep networks, e.g. ResNEXT (Xie et al., 2017) or the mixture of experts (MoE) module adopted in Shazeer et al. (2017). But a major difference here is that all components (layers, channels, and operators) of the ReduNet are by explicit construction from first principles and they all have precise optimization, statistical and geometric interpretation. Furthermore, there is no need to learn them from back-propagation, although in principle one still could if further fine-tuning of the network is needed (see Remark 3 of Appendix A for more discussions).

2.3 DEEP CONVOLUTION NETWORKS FROM SHIFT-INVARIANT RATE REDUCTION

We next examine ReduNet from the perspective of invariance to transformation. Using the basic and important case of shift/translation invariance as an example, we will show that for data which are compatible with an invariant classifier, the ReduNet construction automatically takes the form of a (*multi-channel*) *convolutional neural network*, rather than heuristically imposed upon.

⁴The choice of the softmax is mostly for its simplicity as it is widely used in other (forward components of) deep networks for purposes such as selection, gating (Shazeer et al., 2017) and routing (Sabour et al., 2017). In principle, this term can be approximated by other operators, say using ReLU that is more amenable to training with back propagation, see Remark 3 in Appendix A.

1D Serial Data and Shift Invariance. For one-dimensional data $\mathbf{x} = [x(0), x(1), \dots, x(n-1)] \in \mathbb{R}^n$ under shift symmetry, we take \mathbb{G} to be the group of circular shifts. Each observation \mathbf{x}^i generates a family $\{\mathbf{x}^i \circ \mathbf{g} \mid \mathbf{g} \in \mathbb{G}\}$ of shifted copies, which are the columns of the circulant matrix $\text{circ}(\mathbf{x}^i) \in \mathbb{R}^{n \times n}$ given by

$$\text{circ}(\mathbf{x}) \doteq \begin{bmatrix} x(0) & x(n-1) & \dots & x(2) & x(1) \\ x(1) & x(0) & x(n-1) & \dots & x(2) \\ \vdots & x(1) & x(0) & \ddots & \vdots \\ x(n-2) & \vdots & \ddots & \ddots & x(n-1) \\ x(n-1) & x(n-2) & \dots & x(1) & x(0) \end{bmatrix} \in \mathbb{R}^{n \times n}.$$

We refer the reader to Appendix B.1 or Kra & Simanca (2012) for properties of circulant matrices.

What happens if we construct the ReduNet from these families $\mathbf{Z}_1 = [\text{circ}(\mathbf{x}^1), \dots, \text{circ}(\mathbf{x}^m)]$? The data covariance matrix:

$$\mathbf{Z}_1 \mathbf{Z}_1^* = [\text{circ}(\mathbf{x}^1), \dots, \text{circ}(\mathbf{x}^m)] [\text{circ}(\mathbf{x}^1), \dots, \text{circ}(\mathbf{x}^m)]^* = \sum_{i=1}^m \text{circ}(\mathbf{x}^i) \text{circ}(\mathbf{x}^i)^* \in \mathbb{R}^{n \times n}$$

associated with this family of samples is *automatically* a (symmetric) circulant matrix. Moreover, because the circulant property is preserved under sums, inverses, and products, the matrices \mathbf{E}_1 and \mathbf{C}_1^j are also automatically circulant matrices, whose application to a feature vector $\mathbf{z} \in \mathbb{R}^n$ can be implemented using circular convolution “ \otimes ”. Specifically, we have the following proposition.

Proposition 2.1 (Convolution structures of \mathbf{E}_1 and \mathbf{C}_1^j) *The matrix $\mathbf{E}_1 = \alpha(\mathbf{I} + \alpha \mathbf{Z}_1 \mathbf{Z}_1^*)^{-1}$ is a circulant matrix and represents a circular convolution:*

$$\mathbf{E}_1 \mathbf{z} = \mathbf{e}_1 \otimes \mathbf{z},$$

where $\mathbf{e}_1 \in \mathbb{R}^n$ is the first column vector of \mathbf{E}_1 and “ \otimes ” is circular convolution defined as

$$(\mathbf{e}_1 \otimes \mathbf{z})_i \doteq \sum_{j=0}^{n-1} e_1(j) x(i + n - j \bmod n).$$

Similarly, the matrices \mathbf{C}_1^j associated with any subsets of \mathbf{Z}_1 are also circular convolutions.

From Proposition 2.1, we have

$$\mathbf{z}_2 \propto \mathbf{z}_1 + \eta \cdot g(\mathbf{z}_1, \boldsymbol{\theta}_1) = \mathbf{z}_1 + \eta \cdot \mathbf{e}_1 \otimes \mathbf{z}_1 - \eta \cdot \boldsymbol{\sigma}([c_1^1 \otimes \mathbf{z}_1, \dots, c_1^k \otimes \mathbf{z}_1]). \quad (15)$$

Because $g(\cdot, \boldsymbol{\theta}_1)$ consists only of operations that co-vary with cyclic shifts, the features \mathbf{Z}_2 at the next level again consist of families of shifts: $\mathbf{Z}_2 = [\text{circ}(\mathbf{x}^1 + \eta g(\mathbf{x}^1, \boldsymbol{\theta}_1)), \dots, \text{circ}(\mathbf{x}^m + \eta g(\mathbf{x}^m, \boldsymbol{\theta}_m))]$. Continuing inductively, we see that all matrices \mathbf{E}_ℓ and \mathbf{C}_ℓ^j based on such \mathbf{Z}_ℓ are circulant. By virtue of the properties of the data, ReduNet has taken the form of a convolutional network, *with no need to explicitly choose this structure!*

The Role of Multiple Channel Lifting and Sparsity. There is one problem though: In general, the set of all circular permutations of a vector \mathbf{z} give a full-rank matrix. That is, the n “augmented” features associated with each sample (hence each class) typically already span the entire space \mathbb{R}^n . The MCR² objective (1) will not be able to distinguish classes as different subspaces.⁵

One natural remedy is to improve the separability of the data by “lifting” the original signal to a higher dimensional space, e.g., by taking their responses to multiple, filters $\mathbf{k}_1, \dots, \mathbf{k}_C \in \mathbb{R}^n$:

$$\mathbf{z}[c] = \mathbf{k}_c \otimes \mathbf{x} = \text{circ}(\mathbf{k}_c) \mathbf{x} \in \mathbb{R}^n, \quad c = 1, \dots, C. \quad (16)$$

The filters can be pre-designed invariance-promoting filters,⁶ or adaptively learned from the data,⁷ or randomly selected as we do in our experiments. This operation lifts each original signal $\mathbf{x} \in \mathbb{R}^n$

⁵All shifted versions delta function $\delta(n)$ can generate any other signal as their (dense) weighted sum.

⁶For 1D signals like audio, one may consider the conventional short time Fourier transform (STFT); for 2D images, one may consider 2D wavelets as in the ScatteringNet (Bruna & Mallat, 2013).

⁷For learned filters, one can learn filters as the principal components of samples as in the PCANet (Chen et al., 2015) or from convolution dictionary learning (Li & Bresler, 2019; Qu et al., 2019).

to a C -channel feature, denoted as $\bar{z} \doteq [z[1], \dots, z[C]]^* \in \mathbb{R}^{C \times n}$. Then, we may construct the ReduNet on vector representations of \bar{z} , denoted as $\text{vec}(\bar{z}) \doteq [z[1]^*, \dots, z[C]^*] \in \mathbb{R}^{nC}$. The associated circulant version $\text{circ}(\bar{z})$ and its data covariance matrix, denoted as $\bar{\Sigma}$, for all its shifted versions are given as:

$$\text{circ}(\bar{z}) \doteq \begin{bmatrix} \text{circ}(z[1]) \\ \vdots \\ \text{circ}(z[C]) \end{bmatrix} \in \mathbb{R}^{nC \times n}, \quad \bar{\Sigma} \doteq \begin{bmatrix} \text{circ}(z[1]) \\ \vdots \\ \text{circ}(z[C]) \end{bmatrix} [\text{circ}(z[1])^*, \dots, \text{circ}(z[C])^*] \in \mathbb{R}^{nC \times nC}, \quad (17)$$

where $\text{circ}(z[c]) \in \mathbb{R}^{n \times n}$ with $c \in [C]$ is the circulant version of the c -th channel of the feature \bar{z} . Then the columns of $\text{circ}(\bar{z})$ will only span at most an n -dimensional proper subspace in \mathbb{R}^{nC} .

However, this simple (linear) lifting operation is not sufficient to render the classes separable yet – features associated with other classes will span the *same* n -dimensional subspace. This reflects a fundamental conflict between linear (subspace) modeling and invariance.

One way of resolving this conflict is to leverage additional structure within each class, in the form of *sparsity*: Signals within each class are not generated as arbitrary linear combinations of some base atoms (or motifs), but only *sparse* combinations of them and their shifted versions. Let \mathcal{D}_j denote a collection of atoms associated for class j , also known as a dictionary, then each signal x in this class is sparsely generated as:

$$x = \text{circ}(\mathcal{D}_j)z$$

for some sparse vector z . Signals in different classes are then generated by different dictionaries whose atoms (or motifs) are incoherent from one another. Due to incoherence, signals in one class are unlikely to be sparsely represented by atoms in any other class. Hence all signals in the k class can be represented as

$$x = [\text{circ}(\mathcal{D}_1), \text{circ}(\mathcal{D}_2), \dots, \text{circ}(\mathcal{D}_k)]\bar{z}$$

where \bar{z} is sparse.⁸ There is a vast literature on how to learn the most compact and optimal sparsifying dictionaries from sample data, e.g. (Li & Bresler, 2019; Qu et al., 2019) and subsequently solve the inverse problem and compute the associated sparse code z or \bar{z} .

Nevertheless, here we are not interested in the optimal dictionary and the precise sparse code for each individual signal. We are only interested if the set of sparse codes for each class are collectively separable from those of other classes. Under the assumption of the sparse generative model, if the convolution kernels $\{k_c\}$ match well with the “transpose” or “inverse” of the above sparsifying dictionaries, also known as the *analysis filters* (Nam et al., 2013; Rubinstein & Elad, 2014), signals in one class will only have high responses to a small subset of those filters and low responses to others (due to the incoherence assumption). Nevertheless, in practice, often a sufficient number of random filters suffice the purpose of ensuring features of different classes have different response patterns to different filters hence make different classes separable (Chan et al., 2015). As optimal sparse coding is not the focus of this paper, we will use the simple random filter design in our experiments, which is adequate to verify the concept.⁹

Hence the multi-channel responses \bar{z} should be sparse. So to approximate the sparse code \bar{z} , we may take an entry-wise *sparsity-promoting nonlinear thresholding*, say $\tau(\cdot)$, on the filter outputs by setting low (say absolute value below ϵ) or negative responses to be zero:

$$\bar{z} = \tau[\text{circ}(k_1)x, \dots, \text{circ}(k_C)x] \in \mathbb{R}^{n \times C}.$$

One may refer to (Rubinstein & Elad, 2014) for a more systematical study on the design of the sparsifying thresholding operator. Nevertheless, here we are not so interested in obtaining the best sparse codes as long as the codes are sufficiently separable. Hence the nonlinear operator τ can be simply chosen to be a soft thresholding or a ReLU. These presumably sparse features \bar{z} can be assumed to lie on a lower-dimensional (nonlinear) submanifold of $\mathbb{R}^{n \times C}$, which can be linearized and separated from the other classes by subsequent ReduNet layers, as illustrated in Figure 3.

⁸Notice that similar sparse representation models have long been proposed and used for classification purposes in applications such a face recognition, demonstrating excellent effectiveness (Wright et al., 2009; Wagner et al., 2012).

⁹Although better sparse coding schemes may surely lead to better classification performance, at a higher computational cost.

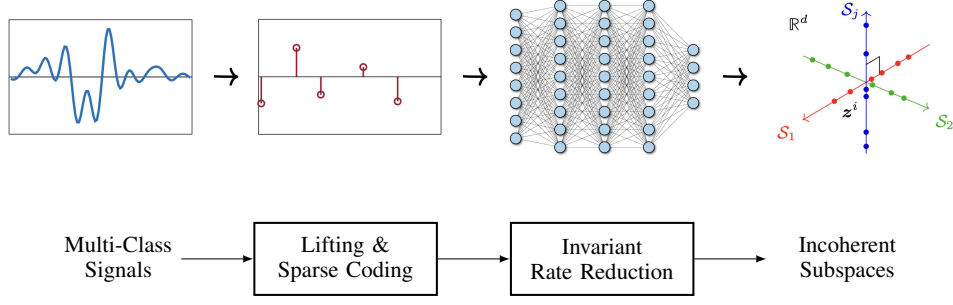


Figure 3: Overview of the process for classifying multi-class signals with shift invariance: Multi-channel lifting and sparse coding followed by a (convolutional) ReduNet for invariant rate reduction. These operations are *necessary* to map shift-invariant multi-class signals to incoherent (linear) subspaces. Note that most modern deep neural networks resemble this process.

The ReduNet constructed from circulant version of these multi-channel features $\bar{\mathbf{z}}$, i.e., $\text{circ}(\bar{\mathbf{Z}}) \doteq [\text{circ}(\bar{\mathbf{z}}^1), \dots, \text{circ}(\bar{\mathbf{z}}^n)]$, retains the good invariance properties described above: the linear operators, now denoted as $\bar{\mathbf{E}}$ and $\bar{\mathbf{C}}^j$, remain block circulant, and represent *multi-channel 1D circular convolutions*. Specifically, we have the following result (see Appendix B.2 for a proof).

Proposition 2.2 (Multi-channel convolution structures of $\bar{\mathbf{E}}$ and $\bar{\mathbf{C}}^j$) *The matrix*

$$\bar{\mathbf{E}} \doteq \alpha (\mathbf{I} + \alpha \text{circ}(\bar{\mathbf{Z}}) \text{circ}(\bar{\mathbf{Z}})^*)^{-1} \quad (18)$$

is block circulant, i.e.,

$$\bar{\mathbf{E}} = \begin{bmatrix} \bar{\mathbf{E}}_{1,1} & \cdots & \bar{\mathbf{E}}_{1,C} \\ \vdots & \ddots & \vdots \\ \bar{\mathbf{E}}_{C,1} & \cdots & \bar{\mathbf{E}}_{C,C} \end{bmatrix} \in \mathbb{R}^{nC \times nC},$$

where each $\bar{\mathbf{E}}_{c,c'} \in \mathbb{R}^{n \times n}$ is a circulant matrix. Moreover, $\bar{\mathbf{E}}$ represents a multi-channel circular convolution, i.e., for any multi-channel signal $\bar{\mathbf{z}} \in \mathbb{R}^{C \times n}$ we have

$$\bar{\mathbf{E}} \cdot \text{vec}(\bar{\mathbf{z}}) = \text{vec}(\bar{\mathbf{e}} \circledast \bar{\mathbf{z}}).$$

In above, $\bar{\mathbf{e}} \in \mathbb{R}^{C \times C \times n}$ is a multi-channel convolutional kernel with $\bar{\mathbf{e}}[c, c'] \in \mathbb{R}^n$ being the first column vector of $\bar{\mathbf{E}}_{c,c'}$, and $\bar{\mathbf{e}} \circledast \bar{\mathbf{z}} \in \mathbb{R}^{C \times n}$ is the multi-channel circular convolution defined as

$$(\bar{\mathbf{e}} \circledast \bar{\mathbf{z}})[c] \doteq \sum_{c'=1}^C \bar{\mathbf{e}}[c, c'] \circledast \bar{\mathbf{z}}[c'], \quad \forall c = 1, \dots, C.$$

Similarly, the matrices $\bar{\mathbf{C}}^j$ associated with any subsets of $\bar{\mathbf{Z}}$ are also multi-channel circular convolutions.

From Proposition 2.2, ReduNet is a deep convolutional network for multi-channel 1D signals by construction.¹⁰ Figure 3 illustrates the whole process of rate reduction with such sparse and invariant features.

Fast Computation in the Spectral Domain. The calculation of $\bar{\mathbf{E}}$ in (18) requires inverting a matrix of size $nC \times nC$, which has complexity $O(n^3 C^3)$. By using the relationship between circulant matrix and Discrete Fourier Transform (DFT) of a 1D signal, this complexity can be significantly reduced. Specifically, let $\mathbf{F} \in \mathbb{C}^{n \times n}$ be the DFT matrix, and $\text{DFT}(\mathbf{z}) \doteq \mathbf{F}\mathbf{z} \in \mathbb{C}^{n \times n}$ be the DFT of $\mathbf{z} \in \mathbb{R}^n$, where \mathbb{C} denotes the set of complex numbers. We have

$$\text{circ}(\mathbf{z}) = \mathbf{F}^* \text{diag}(\text{DFT}(\mathbf{z})) \mathbf{F}. \quad (19)$$

We refer the reader to Appendix B.3 for properties of circulant matrices and DFT. By using the relation in (19), $\bar{\mathbf{E}}$ can be computed as

$$\bar{\mathbf{E}} = \begin{bmatrix} \mathbf{F}^* & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \ddots & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{F}^* \end{bmatrix} \cdot \alpha \left(\mathbf{I} + \alpha \begin{bmatrix} D_{11} & \cdots & D_{1C} \\ \vdots & \ddots & \vdots \\ D_{C1} & \cdots & D_{CC} \end{bmatrix} \right)^{-1} \cdot \begin{bmatrix} \mathbf{F} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \ddots & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{F} \end{bmatrix},$$

¹⁰Unlike Xception nets (Chollet, 2017), these multi-channel convolutions in general are *not* depthwise separable. It remains open what additional structures on the data would lead to depthwise separable convolutions.

where $\mathbf{D}_{cc'} \doteq \sum_{i=1}^m \text{diag}(\text{DFT}(\mathbf{z}^i[c])) \cdot \text{diag}(\text{DFT}(\mathbf{z}^i[c']))^* \in \mathbb{C}^{n \times n}$ is a diagonal matrix. The matrix in the inverse operator is a block diagonal matrix after a permutation of rows and columns. Hence, to compute $\bar{\mathbf{E}}$ and $\bar{\mathbf{C}}^j \in \mathbb{R}^{nC \times nC}$, we only need to compute in the frequency domain the inverse of $C \times C$ blocks for n times and the overall complexity is $O(nC^3)$.

The benefit of computation with DFT motivates us to construct the ReduNet in the spectral domain. Let $\{\bar{\mathbf{z}}^i \in \mathbb{R}^{n \times C}\}_{i=1}^m$ be a collection of multi-channel 1D signals and represent it as a matrix $\bar{\mathbf{Z}} \in \mathbb{R}^{C \times n \times m}$. Then, we define the *shift invariant coding rate reduction* for $\bar{\mathbf{Z}}$ as

$$\begin{aligned} \Delta R_{\text{circ}}(\bar{\mathbf{Z}}, \mathbf{\Pi}) &\doteq \frac{1}{n} \Delta R(\text{circ}(\bar{\mathbf{Z}}), \bar{\mathbf{\Pi}}) \\ &= \frac{1}{2n} \log \det \left(\mathbf{I} + \alpha \text{circ}(\bar{\mathbf{Z}}) \text{circ}(\bar{\mathbf{Z}})^* \right) - \sum_{j=1}^k \frac{\gamma_j}{2n} \log \det \left(\mathbf{I} + \alpha_j \text{circ}(\bar{\mathbf{Z}}) \bar{\mathbf{\Pi}}^j \text{circ}(\bar{\mathbf{Z}})^* \right), \end{aligned}$$

where $\alpha = \frac{Cn}{mn\epsilon^2} = \frac{C}{m\epsilon^2}$, $\alpha_j = \frac{Cn}{\text{tr}(\mathbf{\Pi}^j)n\epsilon^2} = \frac{C}{\text{tr}(\mathbf{\Pi}^j)\epsilon^2}$, $\gamma_j = \frac{\text{tr}(\mathbf{\Pi}^j)}{m}$, and $\bar{\mathbf{\Pi}}^j$ is augmented membership matrix in an obvious way. The normalization factor n is introduced because the circulant matrix $\text{circ}(\bar{\mathbf{Z}})$ contains n (shifted) copies of each signal.

Let $\text{DFT}(\bar{\mathbf{Z}}) \in \mathbb{C}^{C \times n \times m}$ be data in spectral domain obtained by taking DFT on each channel of each signal $\bar{\mathbf{z}}^i$ and denote $\text{DFT}(\bar{\mathbf{Z}})(p) \in \mathbb{C}^{C \times m}$ the p -th slice of $\text{DFT}(\bar{\mathbf{Z}})$ on the second dimension. Then, the gradient of $\Delta R_{\text{circ}}(\bar{\mathbf{Z}}, \mathbf{\Pi})$ w.r.t. $\bar{\mathbf{Z}}$ can be computed from the expansion $\bar{\mathcal{E}} \in \mathbb{C}^{C \times C \times n}$ and compression $\bar{\mathcal{C}}^j \in \mathbb{C}^{C \times C \times n}$ operators in the spectral domain, defined as

$$\begin{aligned} \bar{\mathcal{E}}(p) &\doteq \alpha \cdot [\mathbf{I} + \alpha \cdot \text{DFT}(\bar{\mathbf{Z}})(p) \cdot \text{DFT}(\bar{\mathbf{Z}})(p)^*]^{-1} \in \mathbb{C}^{C \times C}, \\ \bar{\mathcal{C}}^j(p) &\doteq \alpha_j \cdot [\mathbf{I} + \alpha_j \cdot \text{DFT}(\bar{\mathbf{Z}})(p) \cdot \mathbf{\Pi}_j \cdot \text{DFT}(\bar{\mathbf{Z}})(p)^*]^{-1} \in \mathbb{C}^{C \times C}. \end{aligned}$$

In above, $\bar{\mathcal{E}}(p)$ (resp., $\bar{\mathcal{C}}^j(p)$) is the p -th slice of $\bar{\mathcal{E}}$ (resp., $\bar{\mathcal{C}}^j$) on the last dimension. Specifically, we have the following result (see Appendix B.3 for a proof).

Theorem 2.3 (Computing multi-channel convolutions $\bar{\mathbf{E}}$ and $\bar{\mathbf{C}}^j$) Let $\bar{\mathbf{U}} \in \mathbb{C}^{C \times n \times m}$ and $\bar{\mathbf{W}}^j \in \mathbb{C}^{C \times n \times m}$, $j = 1, \dots, k$ be given by

$$\begin{aligned} \bar{\mathbf{U}}(p) &\doteq \bar{\mathcal{E}}(p) \cdot \text{DFT}(\bar{\mathbf{Z}})(p), \\ \bar{\mathbf{W}}^j(p) &\doteq \bar{\mathcal{C}}^j(p) \cdot \text{DFT}(\bar{\mathbf{Z}})(p), \quad j = 1, \dots, k, \end{aligned}$$

for each $p \in \{0, \dots, n-1\}$. Then, we have

$$\begin{aligned} \frac{1}{2n} \frac{\partial \log \det(\mathbf{I} + \alpha \cdot \text{circ}(\bar{\mathbf{Z}}) \text{circ}(\bar{\mathbf{Z}})^*)}{\partial \bar{\mathbf{Z}}} &= \text{IDFT}(\bar{\mathbf{U}}), \\ \frac{\gamma_j}{2n} \frac{\partial \log \det(\mathbf{I} + \alpha_j \cdot \text{circ}(\bar{\mathbf{Z}}) \bar{\mathbf{\Pi}}^j \text{circ}(\bar{\mathbf{Z}})^*)}{\partial \bar{\mathbf{Z}}} &= \gamma_j \cdot \text{IDFT}(\bar{\mathbf{W}}^j \mathbf{\Pi}^j). \end{aligned}$$

In above, $\text{IDFT}(\bar{\mathbf{U}})$ is the time domain signal obtained by taking inverse DFT on each channel of each signal in $\bar{\mathbf{U}}$.

By this result, the gradient ascent update in (3) (when applied to $\Delta R_{\text{circ}}(\bar{\mathbf{Z}}, \mathbf{\Pi})$) can be equivalently expressed as an update in spectral domain on $\bar{\mathbf{V}}_\ell \doteq \text{DFT}(\bar{\mathbf{Z}}_\ell)$ as

$$\bar{\mathbf{V}}_{\ell+1}(p) \propto \bar{\mathbf{V}}_\ell(p) + \eta \left(\bar{\mathcal{E}}_\ell(p) \cdot \bar{\mathbf{V}}_\ell(p) - \sum_{j=1}^k \gamma_j \bar{\mathcal{C}}_\ell^j(p) \cdot \bar{\mathbf{V}}_\ell(p) \mathbf{\Pi}^j \right), \quad p = 0, \dots, n-1,$$

and a ReduNet can be constructed in a similar fashion as before. For implementation details, we refer the reader to Algorithm 1 of Appendix B.3.

Connections to Recurrent and Convolutional Sparse Coding. The sparse coding perspective of Gregor & LeCun (2010) has later been extended to recurrent and convolutional networks for serial data, e.g. Wisdom et al. (2016); Pappayan et al. (2016); Sulam et al. (2018); Monga et al. (2019). Although both sparsity and convolution have long been advocated as desired characteristics for deep networks, their necessity and precise role have never been clearly and rigorously justified, at least not directly from the objective of the network, say classification. In our framework, we see how multi-channel convolutions ($\bar{\mathbf{E}}, \bar{\mathbf{C}}^j$), different nonlinear activations ($\hat{\pi}^j, \tau$), and the sparsity requirement are derived from, rather than heuristically proposed for, the objective of maximizing rate reduction of the features while enforcing shift invariance.

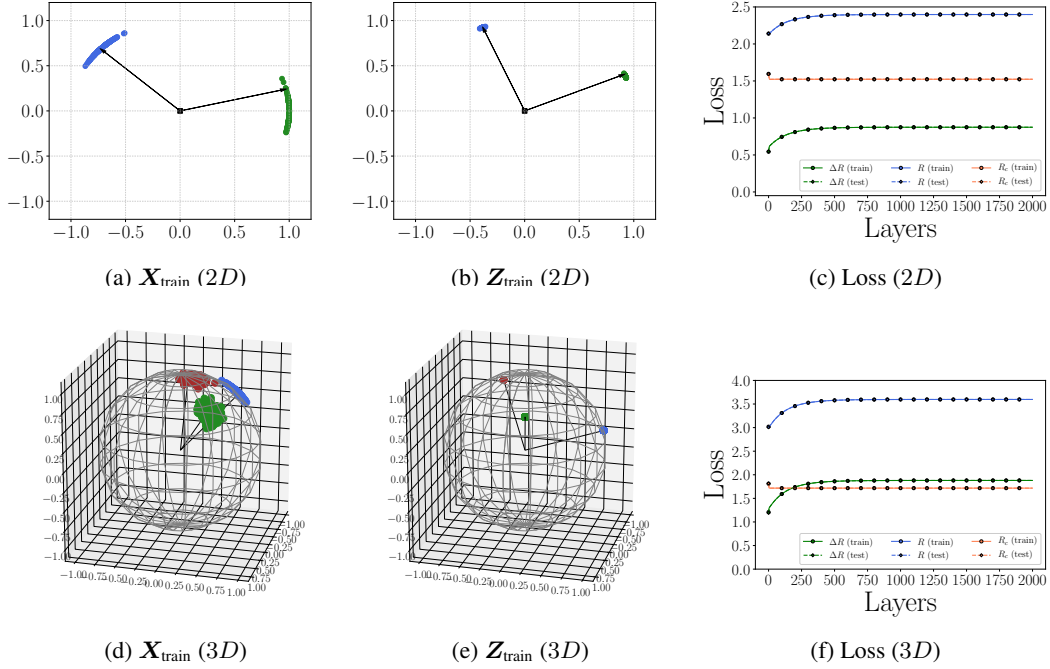


Figure 4: Original samples and learned representations for 2D and 3D Mixture of Gaussians. We visualize data points \mathbf{X} (before mapping) and features \mathbf{Z} (after mapping) by scatter plot. In each scatter plot, each color represents one class of samples. We also show the plots for the progression of values of the objective functions.

2D Images and Translation Invariance. In the case of classifying images invariant to arbitrary 2D translation, for simplicity we may view the image (feature) $\mathbf{z} \in \mathbb{R}^{(W \times H) \times C}$ as a function defined on a torus \mathcal{T}^2 (discretized as a $W \times H$ grid) and consider \mathbb{G} to be the (Abelian) group of all 2D (circular) translations on the torus (see Figure 7 in the experiment section for visualization). As we will show in the Appendix C, the associated linear operators \bar{E} and \bar{C}^j 's act on the image feature \mathbf{z} as *multi-channel 2D circular convolutions*. The resulting network will be a deep convolutional network that shares the same multi-channel convolution structures as conventional CNNs for 2D images (LeCun et al., 1995; Krizhevsky et al., 2012). The difference is that, again, the architectures and parameters of our network are derived from the rate reduction objective, and so are the non-linear activation $\hat{\pi}^j$ and τ . Again, our derivation in Appendix C shows that this multi-channel 2D convolutional network can be constructed more efficiently in the spectral domain (see Theorem C.1 of Appendix C for a rigorous statement and justification).

Sparse Coding and Spectral Computing in Nature. Interestingly, there have been strong scientific evidences that neurons in the visual cortex encode and transmit information in the rate of spiking, hence the so-called “spiking neurons” (Softky & Koch, 1993; Eliasmith & Anderson, 2003; Belitski et al., 2008). Notice that sparse coding is also a main characteristic of the visual cortex (Olshausen & Field, 1996). So remarkably, nature might have already “learned” to exploit benefits of the above mathematical principles, in particular the computational efficiency in sparse coding and in the spectral domain for achieving invariant (visual) recognition!

3 SIMULATIONS AND EXPERIMENTS

We now *verify* whether the so constructed ReduNet achieves its design objectives through experiments on synthetic data and real images. The datasets and experiments are chosen to clearly demonstrate the behaviors of the network obtained by our algorithm, in terms of learning the correct discriminative representation and truly achieving invariance. It is not the purpose of this work to push the state of the art on any real datasets with highly engineered networks and systems, although we believe this framework has this potential in the future. All code is implemented in Python mainly

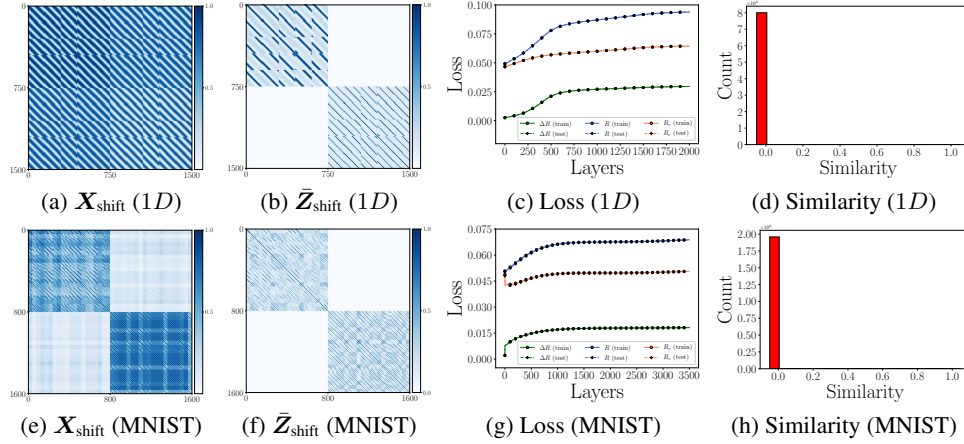


Figure 5: Heatmaps of cosine similarity between data $\mathbf{X}_{\text{shift}}$ /learned features $\bar{\mathbf{Z}}_{\text{shift}}$, MCR² loss, and distance between shift samples and subspaces. For (a), (b), (e), (f), we pick one sample from each class and augment the sample with its every possible shifted ones, then calculate the cosine similarity between these augmented samples. For (d), (h), we first augment each samples in the dataset with its every possible shifted ones, then we evaluate the cosine similarity (in absolute value) between pairs across classes: for each pair, one sample is from training and one sample is from test which belong to different classes.

using NumPy. All our experiments are conducted in a computer with 2.8 GHz Intel i7 CPU and 16GB of memory. Implementation details and more experiments can be found in Appendix D.

Learning Mixture of Gaussians in \mathbb{S}^1 and \mathbb{S}^2 . Consider a mixture of two Gaussian distributions in \mathbb{R}^2 that is projected onto \mathbb{S}^1 . We first generate data points from these two distributions, $\mathbf{X}_1 = [\mathbf{x}_1^1, \dots, \mathbf{x}_1^m] \in \mathbb{R}^{2 \times m}$, $\mathbf{x}_1^i \sim \mathcal{N}(\boldsymbol{\mu}_1, \sigma_1 \mathbf{I})$, and $\pi(\mathbf{x}_1^i) = 1$; $\mathbf{X}_2 = [\mathbf{x}_2^1, \dots, \mathbf{x}_2^m] \in \mathbb{R}^{2 \times m}$, $\mathbf{x}_2^i \sim \mathcal{N}(\boldsymbol{\mu}_2, \sigma_2 \mathbf{I})$, and $\pi(\mathbf{x}_2^i) = 2$. We set $m = 500$, $\sigma_1 = \sigma_2 = 0.1$ and $\boldsymbol{\mu}_1, \boldsymbol{\mu}_2 \in \mathbb{S}^1$. Then we project all the data points onto \mathbb{S}^1 , i.e., $\mathbf{x}_j^i / \|\mathbf{x}_j^i\|_2$. To construct the network (computing \mathbf{E}, \mathbf{C}^j for each layer), we set the # of iterations/layers $L = 2,000$,¹¹ step size $\eta = 0.5$, and precision $\epsilon = 0.1$. As shown in Figure 4a-4b, we can observe that after the mapping $f(\cdot, \boldsymbol{\theta})$, samples from the same class converge to a single cluster and the angle between two different clusters is approximately $\pi/2$, which is well aligned with the optimal solution \mathbf{Z}_* of the MCR² loss in \mathbb{S}^1 . MCR² loss of features on different layers can be found in Figure 4c. Empirically, we find that our constructed network is able to maximize MCR² loss and converges stably. Similarly, we consider mixture of three Gaussian distributions in \mathbb{R}^3 with means $\boldsymbol{\mu}_1, \boldsymbol{\mu}_2, \boldsymbol{\mu}_3$ uniformly in \mathbb{S}^2 , and variance $\sigma_1 = \sigma_2 = \sigma_3 = 0.1$, and all data points are projected onto \mathbb{S}^2 (See Figure 4d-4f). We can observe similar behavior as in \mathbb{S}^2 , i.e., samples from the same class converge to one cluster and different clusters are orthogonal to each other. Moreover, we sample new data points from the same distributions for both cases and find that new samples form the same class consistently converge to the same cluster as the training samples. More examples and details can be found in Appendix D.

Learning Shift Invariant Features. As described in § 2.3, by maximizing the rate reduction via Eq. (15), we are able to explicitly construct operators that are invariant to (circular) shifts. To verify the effectiveness of our proposed network on shift invariance tasks, we apply our network to classify signals sampled from two different 1D functions. The underlying function of the first class is sinusoidal signal $h_1(t) = \sin(t) + \epsilon$, and the second class is a composition of sign and sin function, $h_2(t) = \text{sign}(\sin(t)) + \epsilon$, where $\epsilon \sim \mathcal{N}(0, 0.1)$. (See Figure 12 in Appendix D). Each sample is generated by first picking $t_0 \in [0, 10\pi]$, then obtaining n equidistant point within the boundaries $[t_0, t_0 + 2\pi]$ with i.i.d Gaussian noise. Detailed implementations for sampling from h_1 and h_2 can be found in Appendix D.3. We generate a dataset which contains m samples, with $m/2$ samples in each class, i.e., $\mathbf{X} = [\mathbf{X}_1, \mathbf{X}_2] \in \mathbb{R}^{n \times m}$. Then each sample is lifted to a C -channel feature as defined in (16), i.e., $\bar{\mathbf{X}} \in \mathbb{R}^{(n \cdot C) \times m}$. For training data, We set the number of features $n = 150$, samples $m = 400$, channels $C = 7$, iterations/layers $L = 2,000$, step size $\eta = 0.1$, and precision $\epsilon = 0.1$. We sample the same number of test data points followed by the

¹¹It is remarkable to see how easily our framework leads to working deep networks with thousands of layers! But this also indicates the efficiency of the layers is not so high. Remark 4 provides possible ways to improve.




Figure 6: Examples of rotated images of MNIST digits for testing rotation invariance, each rotated by 18° . **(Left)** Diagram for polar coordinate representation; **(Right)** Rotated digit ‘0’ and digit ‘1’.




Figure 7: Examples of translated images of MNIST digits (with `stride=7`) for testing cyclic translation invariance of the ReduNet. **(Left)** Diagram for cyclic translation; **(Right)** Translated digit ‘0’ and digit ‘1’.

same procedure. As shown in Figure 13, we observe that the network can map the two classes of signals to orthogonal subspaces both on training and test datasets. To verify invariance property of the network, we first pick 5 signal samples from each class (from test dataset) and get their corresponding augmented samples by shifting. Then we have $m = 1,500$ augmented samples for each original signal, $\mathbf{X}_{\text{shift}} \in \mathbb{R}^{150 \times 1,500}$, and we visualize the pairwise inner product of $\mathbf{X}_{\text{shift}}$ and their representations, $\mathbf{Z}_{\text{shift}} \in \mathbb{R}^{(150 \cdot 7) \times 1,500}$ in Figure 5a-5b. Moreover, we augment every sample (from the test dataset) with its all possible shifted versions and calculate the cosine similarity between their representations and the representations of all the training samples from the other class (in Figure 5d). We find that the proposed network can map different classes of signals (including all shifted augmentations) to orthogonal subspaces, to increase the MCR² loss (shown in Figure 5c).

Rotational Invariance on MNIST Digits. We study the ReduNet on learning *rotation* invariant features on MNIST dataset (LeCun, 1998). We impose a polar grid on the image $\mathbf{x} \in \mathbb{R}^{H \times W}$, with its geometric center being the center of the 2D polar grid. For each radius $r_i, i \in [C]$, we can sample Γ pixels with respect to each angle $\gamma_l = l \cdot (2\pi/\Gamma)$ with $l \in [\Gamma]$. Then given an image sample \mathbf{x} from the dataset, we represent the image in a polar coordinate representation $\mathbf{x}(p) = (\gamma_{l,i}, r_{l,i}) \in \mathbb{R}^{\Gamma \times C}$. Examples of rotated images are shown in Figure 6. Our goal is to learn rotation invariant features, i.e., we expect to learn $f(\cdot, \theta)$ such that $\{f(\mathbf{x}(p) \circ \mathbf{g}, \theta)\}_{\mathbf{g} \in \mathbb{G}}$ lie in the same subspace, where \mathbf{g} is the shift transformation in polar angle. By performing polar coordinate transformation for images from digit ‘0’ and digit ‘1’ in the training dataset, we can obtain the data matrix $\mathbf{X}(p) \in \mathbb{R}^{(\Gamma \cdot C) \times m}$. After performing polar coordinate transformation, the rotation operation in the original images is equivalent to the shift operation in the polar coordinate system. We use $m = 2,000$ training samples, set $\Gamma = 200$ and $C = 5$ for polar transformation, and set iteration $L = 3,500$, precision $\epsilon = 0.1$, step-size $\eta = 0.5$. We generate 1,000 test samples followed by the same procedure. In Figure 15, we can see that our proposed ReduNet is able to map most samples from different classes to orthogonal subspaces (w.r.t. class) on test dataset. Meanwhile, in Figure 5e, 5f, and 5h, we observe that the learnt features are invariant to shift transformation in polar angle (i.e., arbitrary rotation in \mathbf{x}).

We compare the accuracy (both on the original test data and the shifted test data) of the ReduNet (without considering invariance) and the shift invariant ReduNet. For ReduNet (without considering invariance), we use the same training dataset as the shift invariant ReduNet, we set iteration $L = 3,500$, step size $\eta = 0.5$, and precision $\epsilon = 0.1$. The results are summarized in Table 1.

Table 1: Comparing network performance on learning rotational-invariant representations and 2D translation-invariant representations on MNIST.

SHIFT INVARIANCE	REDUNET	REDUNET (SHIFT-INVARIANT)
ACC (ORIGINAL TEST DATA)	0.983	0.996
ACC (TEST DATA WITH ALL POSSIBLE SHIFTS)	0.707	0.993
TRANSLATIONAL INVARIANCE	REDUNET	REDUNET (TRANSLATION-INVARIANT)
ACC (ORIGINAL TEST DATA)	0.980	0.975
ACC (TEST DATA WITH ALL POSSIBLE SHIFTS)	0.540	0.909

With the invariant design, we can see from Table 1 that the shift invariant ReduNet achieves better performance in terms of invariance on the MNIST binary classification task.

Translational Invariance on MNIST Digits. We provide experimental results for verifying the invariance property of ReduNet under 2D translations. We construct 1). ReduNet (without considering invariance) and 2). 2D translation-invariant ReduNet for classifying digit ‘0’ and digit ‘1’ on MNIST dataset. We use $m = 1,000$ samples (500 samples from each class) for training the models, and use another 500 samples (250 samples from each class) for evaluation. To evaluate the 2D translational invariance, for each test image $\mathbf{x}_{\text{test}} \in \mathbb{R}^{H \times W}$, we consider *all* translation augmentations of the test image with a `stride=7`. More specifically, for the MNIST dataset, we have $H = W = 28$. So for each image, the total number of all cyclic translation augmentations (with `stride=7`) is $4 \times 4 = 16$. Examples of translated images are shown in Figure 7. Notice that such translations are considerably larger than normally considered in the literature since we consider invariance to the entire group of cyclic translations on the $H \times W$ grid as a torus.

For ReduNet (without considering translation invariance), we set iteration $L = 2,000$, step size $\eta = 0.1$, and precision $\epsilon = 0.1$. For translation-invariant ReduNet, we set $L = 2,000$, step size $\eta = 0.5$, precision $\epsilon = 0.1$, number of channels $C = 5$, and kernel size is set as 3×3 . We summarize the results in Table 1. Similar to the 1D rotational results on the MNIST dataset, the translation-invariant ReduNet achieves better performance under translations compared with the ReduNet without considering invariance. The accuracy drop of the translation-invariant ReduNet is much less than the one of ReduNet without invariance design.

4 CONCLUSIONS AND FUTURE WORK

Following the recently proposed maximal coding rate reduction framework of Yu et al. (2020), this work offers a principled interpretation of modern deep (convolutional) networks by construction from first principles and with minimal assumptions. As we see, the rate reduction principle provides a rigorous explanation for the deep architecture and components from the perspective of optimizing the rate reduction of final representations. In particular, most key characteristics of modern deep neural networks, including their layered architectures, linear (convolutional) operators (for shift invariance), and nonlinear operators (for classifying and sparsifying), can all be derived as necessary operations for optimizing this objective. Somewhat unexpected, our analysis shows that the nominal architecture and parameters of such a deep (convolution) network can all be explicitly constructed layer-by-layer in a forward propagation fashion without the need of back propagation training. The analysis further reveals that such a deep convolutional network is computationally more efficient to construct and learn in the spectral domain. Preliminary simulations and experiments on basic data sets clearly verify the so-constructed ReduNet achieves the desired functionality and objective.

Although in this work the ReduNet is purely forward constructed as the nominal optimization path for rate reduction, one may study how to effectively fine tune it via back propagation. In this work the rate distortion function primarily treats each class as low-dimensional subspace or (degenerate) Gaussian. As already pointed out by Wright et al. (2008), if the data have more sophisticated nonlinear structures, to better classify such data, the rate distortion function $R(\mathbf{Z})$ can readily incorporate an arbitrary kernel function $k(\mathbf{z}, \mathbf{z}')$ by replacing the inner product in $\log \det (\mathbf{I} + \alpha \mathbf{Z} \mathbf{Z}^*) = \log \det (\mathbf{I} + \alpha \mathbf{Z}^* \mathbf{Z})$. Hence this framework is naturally amenable to

further analysis and development with kernel functions. Furthermore, the lossy compression and rate distortion framework was originally developed for seeking optimal data clustering Π via minimizing the second term R_c of ΔR (Ma et al., 2007; Vidal et al., 2016). Hence we believe the maximal rate reduction framework can be naturally extended to settings of online, self-supervised, or *unsupervised* learning if the class information Π is partially or entirely unknown and is to be optimized jointly with the representation Z .

All in all, we strongly believe that maximizing rate reduction provides a principled framework for designing new networks with interpretable architectures and operators that could scale up to real-world datasets and problems, with better performance guarantees.

ACKNOWLEDGEMENTS

Yi acknowledges support from ONR grant N00014-20-1-2002 and the joint Simons Foundation-NSF DMS grant #2031899, as well as support from Berkeley FHL Vive Center for Enhanced Reality and Berkeley Center for Augmented Cognition. Chong and Yi acknowledge support from Tsinghua-Berkeley Shenzhen Institute (TBSI) Research Fund. Yaodong, Haozhi, and Yi acknowledge support from Berkeley AI Research (BAIR). John acknowledges support from NSF grants 1838061, 1740833, and 1733857.

REFERENCES

- Mongi A Abidi, Andrei V Gribok, and Joonki Paik. *Optimization Techniques in Computer Vision*. Springer, 2016.
- Marcin Andrychowicz, Misha Denil, Sergio Gomez, Matthew W Hoffman, David Pfau, Tom Schaul, Brendan Shillingford, and Nando De Freitas. Learning to learn by gradient descent by gradient descent. In *Advances in neural information processing systems*, pp. 3981–3989, 2016.
- Aharon Azulay and Yair Weiss. Why do deep convolutional networks generalize so poorly to small image transformations? *arXiv preprint arXiv:1805.12177*, 2018.
- Bowen Baker, Otkrist Gupta, N. Naik, and R. Raskar. Designing neural network architectures using reinforcement learning. *ArXiv*, abs/1611.02167, 2017.
- Amir Beck and Marc Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202, 2009.
- Andrei Belitski, Arthur Gretton, Cesare Magri, Yusuke Murayama, Marcelo A. Montemurro, Nikos K. Logothetis, and Stefano Panzeri. Low-frequency local field potentials and spikes in primary visual cortex convey independent visual information. *Journal of Neuroscience*, 28(22):5696–5709, 2008. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.0009-08.2008. URL <https://www.jneurosci.org/content/28/22/5696>.
- Joan Bruna and Stéphane Mallat. Invariant scattering convolution networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1872–1886, 2013.
- Tsung-Han Chan, Kui Jia, Shenghua Gao, Jiwen Lu, Zinan Zeng, and Yi Ma. PCANet: A simple deep learning baseline for image classification? *IEEE transactions on image processing*, 24(12):5017–5032, 2015.
- Tian Qi Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. In *Advances in Neural Information Processing Systems*, pp. 6571–6583, 2018.
- François Chollet. Xception: Deep learning with depthwise separable convolutions. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1800–1807, 2017.
- Taco S. Cohen and Max Welling. Group equivariant convolutional networks. *CoRR*, abs/1602.07576, 2016. URL <http://arxiv.org/abs/1602.07576>.
- Dheeru Dua and Casey Graff. Uci machine learning repository, 2017. URL <http://archive.ics.uci.edu/ml>.

- Chris Eliasmith and Charles Anderson. *Neural Engineering: Computation, Representation and Dynamics in Neurobiological Systems*. Cambridge, MA, 01 2003.
- Logan Engstrom, Brandon Tran, Dimitris Tsipras, Ludwig Schmidt, and Aleksander Madry. A rotation and a translation suffice: Fooling CNNs with simple transformations. *arXiv preprint arXiv:1712.02779*, 2017.
- Karol Gregor and Yann LeCun. Learning fast approximations of sparse coding. In *Proceedings of the 27th International Conference on International Conference on Machine Learning*, pp. 399–406, 2010.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Geoffrey E. Hinton, A. Krizhevsky, and S. Wang. Transforming auto-encoders. In *ICANN*, 2011.
- Sepp Hochreiter and Jrgen Schmidhuber. Long short-term memory. *Neural computation*, 9:1735–80, 12 1997. doi: 10.1162/neco.1997.9.8.1735.
- Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2261–2269, 2017.
- Frank Hutter, Lars Kotthoff, and Joaquin Vanschoren (eds.). *Automatic Machine Learning: Methods, Systems, Challenges*. Springer, 2019.
- Chi Jin, Praneeth Netrapalli, and Michael I. Jordan. Accelerated gradient descent escapes saddle points faster than gradient descent. In Sébastien Bubeck, Vianney Perchet, and Philippe Rigollet (eds.), *Conference On Learning Theory, COLT 2018, Stockholm, Sweden, 6-9 July 2018*, volume 75 of *Proceedings of Machine Learning Research*, pp. 1042–1085. PMLR, 2018. URL <http://proceedings.mlr.press/v75/jin18a.html>.
- Irwin Kra and Santiago R Simanca. On circulant matrices. *Notices of the American Mathematical Society*, 59:368–377, 2012.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pp. 1097–1105, 2012.
- Andrew Lavin and Scott Gray. Fast algorithms for convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4013–4021, 2016.
- Yann LeCun. The MNIST database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998.
- Yann LeCun, L.D. Jackel, Leon Bottou, Corinna Cortes, J. S. Denker, Harris Drucker, I. Guyon, U.A. Muller, Eduard Sackinger, Patrice Simard, and V. Vapnik. *Learning algorithms for classification: A comparison on handwritten digit recognition*, pp. 261–276. World Scientific, 1995.
- Yanjun Li and Yoram Bresler. Multichannel sparse blind deconvolution on the sphere. *IEEE Transactions on Information Theory*, 65(11):7415–7436, 2019.
- Yi Ma, Harm Derksen, Wei Hong, and John Wright. Segmentation of multivariate mixed data via lossy data coding and compression. *IEEE transactions on pattern analysis and machine intelligence*, 29(9):1546–1562, 2007.
- Haggai Maron, O. Litany, Gal Chechik, and Ethan Fetaya. On learning sets of symmetric elements. *ArXiv*, abs/2002.08599, 2020.
- Michael Mathieu, Mikael Henaff, and Yann LeCun. Fast training of convolutional networks through ffts, 2013.
- Vishal Monga, Yuelong Li, and Yonina C Eldar. Algorithm unrolling: Interpretable, efficient deep learning for signal and image processing. *arXiv preprint arXiv:1912.10557*, 2019.

- S. Nam, M.E. Davies, M. Elad, and R. Gribonval. The cospase analysis model and algorithms. *Applied and Computational Harmonic Analysis*, 34(1):30 – 56, 2013. ISSN 1063-5203. doi: <https://doi.org/10.1016/j.acha.2012.03.006>. URL <http://www.sciencedirect.com/science/article/pii/S1063520312000450>.
- Yurii Nesterov. A method for unconstrained convex minimization problem with the rate of convergence $O(1/k^2)$. In *Doklady AN USSR*, volume 269, pp. 543–547, 1983.
- Bruno A Olshausen and David J Field. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381(6583):607, 1996.
- Vardan Papyan, Yaniv Romano, and Michael Elad. Convolutional neural networks analyzed via convolutional sparse coding. *Journal of Machine Learning Research*, 18, 07 2016.
- Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in python. *the Journal of machine Learning research*, 12:2825–2830, 2011.
- Qing Qu, Xiao Li, and Zhihui Zhu. A nonconvex approach for exact and efficient multichannel sparse blind deconvolution. In *Advances in Neural Information Processing Systems*, pp. 4017–4028, 2019.
- R. Rubinstein and M. Elad. Dictionary learning for analysis-synthesis thresholding. *IEEE Transactions on Signal Processing*, 62(22):5962–5972, 2014.
- Sara Sabour, Nicholas Frosst, and Geoffrey E. Hinton. Dynamic routing between capsules. *CoRR*, abs/1710.09829, 2017. URL <http://arxiv.org/abs/1710.09829>.
- Hanie Sedghi, Vineet Gupta, and Philip M Long. The singular values of convolutional layers. *arXiv preprint arXiv:1805.10408*, 2018.
- Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarczyk, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. In *ICLR*, 2017. URL <https://openreview.net/pdf?id=BlckMDqlg>.
- Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015.
- William R Softky and Christof Koch. The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *Journal of Neuroscience*, 13(1):334–350, 1993.
- Jeremias Sulam, Vardan Papyan, Yaniv Romano, and Michael Elad. Multilayer convolutional sparse modeling: Pursuit and dictionary learning. *IEEE Transactions on Signal Processing*, 66(15): 4090–4104, 2018.
- Xiaoxia Sun, Nasser M Nasrabadi, and Trac D Tran. Supervised deep sparse coding networks for image classification. *IEEE Transactions on Image Processing*, 29:405–418, 2020.
- Nicolas Vasilache, J. Johnson, Michaël Mathieu, Soumith Chintala, Serkan Piantino, and Y. LeCun. Fast convolutional nets with fbfft: A gpu performance evaluation. *CoRR*, abs/1412.7580, 2015.
- Rene Vidal, Yi Ma, and S. S. Sastry. *Generalized Principal Component Analysis*. Springer Publishing Company, Incorporated, 1st edition, 2016. ISBN 0387878106.
- Andrew Wagner, John Wright, Arvind Ganesh, Zihan Zhou, Hossein Mobahi, and Yi Ma. Toward a practical face recognition system: Robust alignment and illumination by sparse representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(2):372–386, 2012.
- Scott Wisdom, Thomas Powers, James Pitton, and Les Atlas. Interpretable recurrent neural networks using sequential sparse recovery. *ArXiv*, abs/1611.07252, 2016.
- John Wright, Yangyu Tao, Zhouchen Lin, Yi Ma, and Heung-Yeung Shum. Classification via minimum incremental coding length (micl). pp. 1633–1640, 2008.

- John Wright, Allen Y. Yang, Arvind Ganesh, S. Shankar Sastry, and Yi Ma. Robust face recognition via sparse representation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 31(2):210–227, February 2009. ISSN 0162-8828. doi: 10.1109/TPAMI.2008.79. URL <http://dx.doi.org/10.1109/TPAMI.2008.79>.
- Denny Wu and J. Xu. On the optimal weighted ℓ_2 regularization in overparameterized linear regression. *ArXiv*, abs/2006.05800, 2020.
- Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5987–5995, 2017.
- Zitong Yang, Yaodong Yu, Chong You, Jacob Steinhardt, and Yi Ma. Rethinking bias-variance trade-off for generalization of neural networks. In *International Conference on Machine Learning (ICML)*, 2020.
- Yaodong Yu, Kwan Ho Ryan Chan, Chong You, Chaobing Song, and Yi Ma. Learning diverse and discriminative representations via the principle of maximal coding rate reduction. In *Advances in Neural Information Processing Systems*, 2020.
- Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Russ R Salakhutdinov, and Alexander J Smola. Deep sets. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information Processing Systems 30*, pp. 3391–3401. Curran Associates, Inc., 2017. URL <http://papers.nips.cc/paper/6931-deep-sets.pdf>.
- Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. 2017.
- Barret Zoph and Quoc V. Le. Neural architecture search with reinforcement learning. 2017. URL <https://arxiv.org/abs/1611.01578>.

A ADDITIONAL REMARKS AND EXTENSIONS

Remark 3 (Approximate with a ReLU Network) *In practice, there are many other simpler non-linear activation functions that one can use to approximate the membership $\hat{\pi}(\cdot)$ and subsequently the nonlinear operation σ in (12). Notice that the geometric meaning of σ in (12) is to compute the “residual” of each feature against the subspace to which it belongs. So when we restrict all our features to be in the first (positive) quadrant of the feature space,¹² one may approximate this residual using the rectified linear units operation, ReLUs, on $\mathbf{p}_j = \mathbf{C}_\ell^j \mathbf{z}_\ell$ or its orthogonal complement:*

$$\sigma(\mathbf{z}_\ell) \propto \mathbf{z}_\ell - \sum_{j=1}^k \text{ReLU}(\mathbf{P}_\ell^j \mathbf{z}_\ell), \quad (20)$$

where $\mathbf{P}_\ell^j = (\mathbf{C}_\ell^j)^\perp$ is the projection onto the j -th class¹³ and $\text{ReLU}(x) = \max(0, x)$. The above approximation is good under the more restrictive assumption that projection of \mathbf{z}_ℓ on the correct class via \mathbf{P}_ℓ^j is mostly large and positive and yet small or negative for other classes.

The resulting ReduNet will be a network primarily involving ReLU operations and feature normalization (onto \mathbb{S}^{n-1}) between each layer. Although in this work, we have argued that the forward-constructed ReduNet network already works to a large extent, in practice one certainly can conduct back-propagation to further fine tune the so-obtained network, say to correct some remaining errors in predicting labels of the training data. Empirically, people have found that deep networks with ReLU activations are easier to train via back propagation (Krizhevsky et al., 2012).

Remark 4 (Accelerated Optimization via Additional Skip Connections) *Empirically, people have found that additional skip connections across multiple layers may improve the network performance, e.g. the DenseNet (Huang et al., 2017). In our framework, the role of each layer is precisely interpreted as one iterative gradient ascent step for the objective function ΔR . In our experiments (see Section 3), we have observed that the basic gradient scheme sometimes converges slowly, resulting in deep networks with thousands of layers (iterations)! To improve the efficiency of the basic ReduNet, one may consider in the future accelerated gradient methods such as the Nesterov acceleration (Nesterov, 1983) or perturbed accelerated gradient descent (Jin et al., 2018). Say to minimize or maximize a function $h(\mathbf{z})$, such accelerated methods usually take the form:*

$$\begin{cases} \mathbf{p}_{\ell+1} &= \mathbf{z}_\ell + \beta_\ell \cdot (\mathbf{z}_\ell - \mathbf{z}_{\ell-1}), \\ \mathbf{z}_{\ell+1} &= \mathbf{p}_{\ell+1} + \eta \cdot \nabla h(\mathbf{p}_{\ell+1}). \end{cases} \quad (21)$$

Hence they require introducing additional skip connections among three layers $\ell-1$, ℓ and $\ell+1$. For typical convex or nonconvex programs, the above accelerated schemes can often reduce the number of iterations by a magnitude.

¹²Most current neural networks seem to adopt this regime.

¹³ \mathbf{P}_ℓ^j can be viewed as the orthogonal complement to \mathbf{C}_ℓ^j .

B 1D CIRCULAR SHIFT INVARIANCE

It has been long known that to implement a convolutional neural network, one can achieve higher computational efficiency by implementing the network in the spectral domain via the fast Fourier transform (Mathieu et al., 2013; Lavin & Gray, 2016; Vasilache et al., 2015). However, our purpose here is different: We want to show that the linear operators E and C^j derived from the gradient flow of MCR² are naturally convolutions when we enforce shift-invariance rigorously. Their convolution structure is derived from the rate reduction objective, rather than heuristically imposed upon the network. Furthermore, the computation involved in constructing these linear operators has a naturally efficient implementation in the spectral domain via fast Fourier transform. Arguably this work is the first to show multi-channel convolutions, together with other convolution-preserving nonlinear operations in the ReduNet, are both necessary and sufficient to ensure shift invariance.

To be somewhat self-contained and self-consistent, in this section, we first introduce our notation and review some of the key properties of circulant matrices which will be used to characterize the properties of the linear operators E and C^j and to compute them efficiently. The reader may refer to Kra & Simanca (2012) for a more rigorous exposition on circulant matrices.

B.1 PROPERTIES OF CIRCULANT MATRIX AND CIRCULAR CONVOLUTION

Given a vector $\mathbf{z} = [z_0, z_1, \dots, z_{n-1}]^* \in \mathbb{R}^n$,¹⁴ we may arrange all its circular shifted versions in a circulant matrix form as

$$\text{circ}(\mathbf{z}) \doteq \begin{bmatrix} z_0 & z_{n-1} & \dots & z_2 & z_1 \\ z_1 & z_0 & z_{n-1} & \dots & z_2 \\ \vdots & z_1 & z_0 & \ddots & \vdots \\ z_{n-2} & \vdots & \ddots & \ddots & z_{n-1} \\ z_{n-1} & z_{n-2} & \dots & z_1 & z_0 \end{bmatrix} \in \mathbb{R}^{n \times n}. \quad (22)$$

Fact 1 (Convolution as matrix multiplication via circulant matrix) *The multiplication of a circulant matrix $\text{circ}(\mathbf{z})$ with a vector $\mathbf{x} \in \mathbb{R}^n$ gives a circular (or cyclic) convolution, i.e.,*

$$\text{circ}(\mathbf{z}) \cdot \mathbf{x} = \mathbf{z} \circledast \mathbf{x}, \quad (23)$$

where

$$(\mathbf{z} \circledast \mathbf{x})_i = \sum_{j=0}^{n-1} x_j z_{i+n-j \bmod n}. \quad (24)$$

Fact 2 (Properties of circulant matrices) *Circulant matrices have the following properties:*

- *Transpose of a circulant matrix, say $\text{circ}(\mathbf{z})^*$, is circulant;*
- *Multiplication of two circulant matrices is circulant, for example $\text{circ}(\mathbf{z})\text{circ}(\mathbf{z})^*$;*
- *For a non-singular circulant matrix, its inverse is also circulant (hence representing a circular convolution).*

These properties of circulant matrices are extensively used in this work as for characterizing the convolution structures of the operators E and C^j .

Given a set of vectors $[\mathbf{z}^1, \dots, \mathbf{z}^m] \in \mathbb{R}^{n \times m}$, let $\text{circ}(\mathbf{z}^i) \in \mathbb{R}^{n \times n}$ be the circulant matrix for \mathbf{z}^i . Then we have the following:

Proposition B.1 (Convolution structures of E and C^j) *Given a set of vectors $\mathbf{Z} = [\mathbf{z}^1, \dots, \mathbf{z}^m]$, the matrix:*

$$\mathbf{E} = \alpha \left(\mathbf{I} + \alpha \sum_{i=1}^m \text{circ}(\mathbf{z}^i) \text{circ}(\mathbf{z}^i)^* \right)^{-1}$$

¹⁴We use superscript $*$ to indicate (conjugate) transpose of a vector or a matrix

is a circulant matrix and represents a circular convolution:

$$\mathbf{E}\mathbf{z} = \mathbf{e} \circledast \mathbf{z},$$

where \mathbf{e} is the first column vector of \mathbf{E} . Similarly, the matrices \mathbf{C}^j associated with any subsets of \mathbf{Z} are also circular convolutions.

B.2 CIRCULANT MATRIX AND CIRCULANT CONVOLUTION FOR MULTI-CHANNEL SIGNALS

In the remainder of this section, we view \mathbf{z} as a 1D signal such as an audio signal. Since we will deal with the more general case of multi-channel signals, we will use the traditional notation T to denote the temporal length of the signal and C for the number of channels. Conceptually, the “dimension” n of such a multi-channel signal, if viewed as a vector, should be $n = CT$.¹⁵ As we will also reveal additional interesting structures of the operators \mathbf{E} and \mathbf{C}^j in the spectral domain, we use t as the index for time, p for the index of frequency, and c for the index of channel.

Given a multi-channel 1D signal $\bar{\mathbf{z}} \in \mathbb{R}^{C \times T}$, we denote

$$\bar{\mathbf{z}} = \begin{bmatrix} \bar{\mathbf{z}}[1]^* \\ \vdots \\ \bar{\mathbf{z}}[C]^* \end{bmatrix} = [\bar{\mathbf{z}}(0), \bar{\mathbf{z}}(1), \dots, \bar{\mathbf{z}}(T-1)] = \{\bar{\mathbf{z}}[c](t)\}_{c=1, t=0}^{c=C, t=T-1}. \quad (25)$$

To compute the coding rate reduction for a collection of such multi-channel 1D signals, we may flatten the matrix representation into a vector representation by stacking the multiple channels of $\bar{\mathbf{z}}$ as a column vector. In particular, we let

$$\text{vec}(\bar{\mathbf{z}}) = [\bar{\mathbf{z}}[1](0), \bar{\mathbf{z}}1, \dots, \bar{\mathbf{z}}[1](T-1), \bar{\mathbf{z}}[2](0), \dots] \in \mathbb{R}^{(C \times T)}. \quad (26)$$

Furthermore, to obtain shift invariance for the coding rate reduction, we may generate a collection of shifted copies of $\bar{\mathbf{z}}$ (along the temporal dimension). Stacking the vector representations for such shifted copies as column vectors, we obtain

$$\text{circ}(\bar{\mathbf{z}}) \doteq \begin{bmatrix} \text{circ}(\bar{\mathbf{z}}[1]) \\ \vdots \\ \text{circ}(\bar{\mathbf{z}}[C]) \end{bmatrix} \in \mathbb{R}^{(C \times T) \times T}. \quad (27)$$

In above, we overload the notation “ $\text{circ}(\cdot)$ ” defined in (22).

We now consider a collection of m multi-channel 1D signals $\{\bar{\mathbf{z}}^i \in \mathbb{R}^{C \times T}\}_{i=1}^m$. Compactly representing the data by $\bar{\mathbf{Z}} \in \mathbb{R}^{C \times T \times m}$ in which the i -th slice on the last dimension is $\bar{\mathbf{z}}^i$, we denote

$$\bar{\mathbf{Z}}[c] = [\bar{\mathbf{z}}^1[c], \dots, \bar{\mathbf{z}}^m[c]] \in \mathbb{R}^{T \times m}, \quad \bar{\mathbf{Z}}(t) = [\bar{\mathbf{z}}^1(t), \dots, \bar{\mathbf{z}}^m(t)] \in \mathbb{R}^{C \times m}. \quad (28)$$

In addition, we denote

$$\begin{aligned} \text{vec}(\bar{\mathbf{Z}}) &= [\text{vec}(\bar{\mathbf{z}}^1), \dots, \text{vec}(\bar{\mathbf{z}}^m)] \in \mathbb{R}^{(C \times T) \times m}, \\ \text{circ}(\bar{\mathbf{Z}}) &= [\text{circ}(\bar{\mathbf{z}}^1), \dots, \text{circ}(\bar{\mathbf{z}}^m)] \in \mathbb{R}^{(C \times T) \times (T \times m)}. \end{aligned} \quad (29)$$

Then, we define the *shift invariant coding rate reduction* for $\bar{\mathbf{Z}} \in \mathbb{R}^{C \times T \times m}$ as

$$\begin{aligned} \Delta R_{\text{circ}}(\bar{\mathbf{Z}}, \mathbf{\Pi}) &\doteq \frac{1}{T} \Delta R(\text{circ}(\bar{\mathbf{Z}}), \bar{\mathbf{\Pi}}) \\ &= \frac{1}{2T} \log \det \left(\mathbf{I} + \alpha \cdot \text{circ}(\bar{\mathbf{Z}}) \cdot \text{circ}(\bar{\mathbf{Z}})^* \right) - \sum_{j=1}^k \frac{\gamma_j}{2T} \log \det \left(\mathbf{I} + \alpha_j \cdot \text{circ}(\bar{\mathbf{Z}}) \cdot \bar{\mathbf{\Pi}}^j \cdot \text{circ}(\bar{\mathbf{Z}})^* \right), \end{aligned} \quad (30)$$

where $\alpha = \frac{CT}{mT\epsilon^2} = \frac{C}{m\epsilon^2}$, $\alpha_j = \frac{CT}{\text{tr}(\mathbf{\Pi}^j)T\epsilon^2} = \frac{C}{\text{tr}(\mathbf{\Pi}^j)\epsilon^2}$, $\gamma_j = \frac{\text{tr}(\mathbf{\Pi}^j)}{m}$, and $\bar{\mathbf{\Pi}}^j$ is augmented membership matrix in an obvious way. Note that we introduce the normalization factor T in (30) because the circulant matrix $\text{circ}(\bar{\mathbf{Z}})$ contains T (shifted) copies of each signal.

¹⁵Notice that in the main paper, for simplicity, we have used n to indicate both the 1D “temporal” or 2D “spatial” dimension of a signal, just to be consistent with the vector case, which corresponds to T here. All notation should be clear within the context.

By applying (4) and (5), we obtain the derivative of $\Delta R_{\text{circ}}(\bar{\mathbf{Z}}, \mathbf{\Pi})$ as

$$\begin{aligned} \frac{1}{2T} \frac{\partial \log \det \left(\mathbf{I} + \alpha \text{circ}(\bar{\mathbf{Z}}) \text{circ}(\bar{\mathbf{Z}})^* \right)}{\partial \text{vec}(\bar{\mathbf{Z}})} &= \frac{1}{2T} \frac{\partial \log \det \left(\mathbf{I} + \alpha \text{circ}(\bar{\mathbf{Z}}) \text{circ}(\bar{\mathbf{Z}})^* \right)}{\partial \text{circ}(\bar{\mathbf{Z}})} \frac{\partial \text{circ}(\bar{\mathbf{Z}})}{\partial \text{vec}(\bar{\mathbf{Z}})} \\ &= \underbrace{\alpha \left(\mathbf{I} + \alpha \text{circ}(\bar{\mathbf{Z}}) \text{circ}(\bar{\mathbf{Z}})^* \right)^{-1}}_{\bar{\mathbf{E}} \in \mathbb{R}^{(C \times T) \times (C \times T)}} \text{vec}(\bar{\mathbf{Z}}), \end{aligned} \quad (31)$$

$$\begin{aligned} \frac{\gamma_j}{2T} \frac{\partial \log \det \left(\mathbf{I} + \alpha_j \text{circ}(\bar{\mathbf{Z}}) \mathbf{\Pi}^j \text{circ}(\bar{\mathbf{Z}})^* \right)}{\partial \text{vec}(\bar{\mathbf{Z}})} &= \gamma_j \underbrace{\alpha_j \left(\mathbf{I} + \alpha_j \text{circ}(\bar{\mathbf{Z}}) \mathbf{\Pi}^j \text{circ}(\bar{\mathbf{Z}})^* \right)^{-1}}_{\bar{\mathbf{C}}^j \in \mathbb{R}^{(C \times T) \times (C \times T)}} \text{vec}(\bar{\mathbf{Z}}) \mathbf{\Pi}^j. \end{aligned} \quad (32)$$

In the following, we show that $\bar{\mathbf{E}} \cdot \text{vec}(\bar{\mathbf{z}})$ represents a multi-channel circular convolution. Note that

$$\bar{\mathbf{E}} = \alpha \begin{bmatrix} \mathbf{I} + \alpha \sum_{i=1}^m \text{circ}(\mathbf{z}^i[1]) \text{circ}(\mathbf{z}^i[1])^* & \cdots & \sum_{i=1}^m \text{circ}(\mathbf{z}^i[1]) \text{circ}(\mathbf{z}^i[C])^* \\ \vdots & \ddots & \vdots \\ \sum_{i=1}^m \text{circ}(\mathbf{z}^i[C]) \text{circ}(\mathbf{z}^i[1])^* & \cdots & \mathbf{I} + \sum_{i=1}^m \alpha \text{circ}(\mathbf{z}^i[C]) \text{circ}(\mathbf{z}^i[C])^* \end{bmatrix}^{-1}. \quad (33)$$

By using Fact 2, the matrix in the inverse above is a *block circulant matrix*, i.e., a block matrix where each block is a circulant matrix. A useful fact about the inverse of such a matrix is the following.

Fact 3 (Inverse of block circulant matrices) *The inverse of a block circulant matrix is a block circulant matrix (with respect to the same block partition).*

The main result of this subsection is the following.

Proposition B.2 (Convolution structures of $\bar{\mathbf{E}}$ and $\bar{\mathbf{C}}^j$) *Given a collection of multi-channel 1D signals $\{\bar{\mathbf{z}}^i \in \mathbb{R}^{C \times T}\}_{i=1}^m$, the matrix $\bar{\mathbf{E}}$ is a block circulant matrix, i.e.,*

$$\bar{\mathbf{E}} \doteq \begin{bmatrix} \bar{\mathbf{E}}_{1,1} & \cdots & \bar{\mathbf{E}}_{1,C} \\ \vdots & \ddots & \vdots \\ \bar{\mathbf{E}}_{C,1} & \cdots & \bar{\mathbf{E}}_{C,C} \end{bmatrix}, \quad (34)$$

where each $\bar{\mathbf{E}}_{c,c'} \in \mathbb{R}^{T \times T}$ is a circulant matrix. Moreover, $\bar{\mathbf{E}}$ represents a multi-channel circular convolution, i.e., for any multi-channel signal $\bar{\mathbf{z}} \in \mathbb{R}^{C \times T}$ we have

$$\bar{\mathbf{E}} \cdot \text{vec}(\bar{\mathbf{z}}) = \text{vec}(\bar{\mathbf{e}} \circledast \bar{\mathbf{z}}).$$

In above, $\bar{\mathbf{e}} \in \mathbb{R}^{C \times C \times T}$ is a multi-channel convolutional kernel with $\bar{\mathbf{e}}[c, c'] \in \mathbb{R}^T$ being the first column vector of $\bar{\mathbf{E}}_{c,c'}$, and $\bar{\mathbf{e}} \circledast \bar{\mathbf{z}} \in \mathbb{R}^{C \times T}$ is the multi-channel circular convolution (with “ \circledast ” overloading the notation from Eq. (24)) defined as

$$(\bar{\mathbf{e}} \circledast \bar{\mathbf{z}})[c] = \sum_{c'=1}^C \bar{\mathbf{e}}[c, c'] \circledast \bar{\mathbf{z}}[c'], \quad \forall c = 1, \dots, C. \quad (35)$$

Similarly, the matrices $\bar{\mathbf{C}}^j$ associated with any subsets of $\bar{\mathbf{Z}}$ are also multi-channel circular convolutions.

Note that the calculation of $\bar{\mathbf{E}}$ in (33) requires inverting a matrix of size $(C \times T) \times (C \times T)$. In the following, we show that this computation can be accelerated by working in the frequency domain.

B.3 FAST COMPUTATION IN SPECTRAL DOMAIN

Circulant matrix and Discrete Fourier Transform. A remarkable property of circulant matrices is that they all share the same set of eigenvectors that form a unitary matrix. We define the matrix:

$$\mathbf{F}_T \doteq \frac{1}{\sqrt{T}} \begin{bmatrix} \omega_T^0 & \omega_T^0 & \cdots & \omega_T^0 & \omega_T^0 \\ \omega_T^0 & \omega_T^1 & \cdots & \omega_T^{T-2} & \omega_T^{T-1} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \omega_T^0 & \omega_T^{T-2} & \cdots & \omega_T^{(T-2)^2} & \omega_T^{(T-2)(T-1)} \\ \omega_T^0 & \omega_T^{T-1} & \cdots & \omega_T^{(T-2)(T-1)} & \omega_T^{(T-1)^2} \end{bmatrix} \in \mathbb{C}^{T \times T}, \quad (36)$$

where $\omega_T \doteq \exp(-\frac{2\pi\sqrt{-1}}{T})$ is the roots of unit (as $\omega_T^T = 1$). The matrix \mathbf{F}_T is a unitary matrix: $\mathbf{F}_T \mathbf{F}_T^* = \mathbf{I}$ and is the well known *Vandermonde matrix*. Multiplying a vector with \mathbf{F}_T is known as the *discrete Fourier transform* (DFT). Be aware that the conventional DFT matrix differs from our definition of \mathbf{F}_T here by a scale: it does not have the $\frac{1}{\sqrt{T}}$ in front. Here for simplicity, we scale it so that \mathbf{F}_T is a unitary matrix and its inverse is simply its conjugate transpose \mathbf{F}_T^* , columns of which represent the eigenvectors of a circulant matrix (Abidi et al., 2016).

Fact 4 (DFT as matrix-vector multiplication) *The DFT of a vector $\mathbf{z} \in \mathbb{R}^T$ can be computed as*

$$\text{DFT}(\mathbf{z}) \doteq \mathbf{F}_T \cdot \mathbf{z} \in \mathbb{C}^T, \quad (37)$$

where

$$\text{DFT}(\mathbf{z})(p) = \frac{1}{\sqrt{T}} \sum_{t=0}^{T-1} z(t) \cdot \omega_T^{p \cdot t}, \quad \forall p = 0, 1, \dots, T-1. \quad (38)$$

The Inverse Discrete Fourier Transform (IDFT) of a signal $\mathbf{v} \in \mathbb{C}^T$ can be computed as

$$\text{IDFT}(\mathbf{v}) \doteq \mathbf{F}_T^* \cdot \mathbf{v} \in \mathbb{C}^T \quad (39)$$

where

$$\text{IDFT}(\mathbf{v})(t) = \frac{1}{\sqrt{T}} \sum_{p=0}^{T-1} v(p) \cdot \omega_T^{-p \cdot t}, \quad \forall t = 0, 1, \dots, T-1. \quad (40)$$

Regarding the relationship between a circulant matrix (convolution) and discrete Fourier transform, we have:

Fact 5 *An $n \times n$ matrix $\mathbf{M} \in \mathbb{C}^{n \times n}$ is a circulant matrix if and only if it is diagonalizable by the unitary matrix \mathbf{F}_n :*

$$\mathbf{F}_n \mathbf{M} \mathbf{F}_n^* = \mathbf{D} \quad \text{or} \quad \mathbf{M} = \mathbf{F}_n^* \mathbf{D} \mathbf{F}_n, \quad (41)$$

where \mathbf{D} is a diagonal matrix of eigenvalues.

Fact 6 (DFT are eigenvalues of the circulant matrix) *Given a vector $\mathbf{z} \in \mathbb{C}^T$, we have*

$$\mathbf{F}_T \cdot \text{circ}(\mathbf{z}) \cdot \mathbf{F}_T^* = \text{diag}(\text{DFT}(\mathbf{z})) \quad \text{or} \quad \text{circ}(\mathbf{z}) = \mathbf{F}_T^* \cdot \text{diag}(\text{DFT}(\mathbf{z})) \cdot \mathbf{F}_T. \quad (42)$$

That is, the eigenvalues of the circulant matrix associated with a vector are given by its DFT.

Fact 7 (Parseval's theorem) *Given any $\mathbf{z} \in \mathbb{C}^T$, we have $\|\mathbf{z}\|_2 = \|\text{DFT}(\mathbf{z})\|_2$. More precisely,*

$$\sum_{t=0}^{T-1} |z[t]|^2 = \sum_{p=0}^{T-1} |\text{DFT}(\mathbf{z})[p]|^2. \quad (43)$$

This property allows us to easily “normalize” features after each layer onto the sphere \mathbb{S}^{n-1} directly in the spectral domain (see Eq. (13) and (60)).

Circulant matrix and Discrete Fourier Transform for multi-channel signals. We now consider multi-channel 1D signals $\bar{\mathbf{z}} \in \mathbb{R}^{C \times T}$. Let $\text{DFT}(\bar{\mathbf{z}}) \in \mathbb{C}^{C \times T}$ be a matrix where the c -th row is the DFT of the corresponding signal $\mathbf{z}[c]$, i.e.,

$$\text{DFT}(\bar{\mathbf{z}}) \doteq \begin{bmatrix} \text{DFT}(\mathbf{z}[1])^* \\ \vdots \\ \text{DFT}(\mathbf{z}[C])^* \end{bmatrix} \in \mathbb{C}^{C \times T}. \quad (44)$$

Similar to the notation in (25), we denote

$$\begin{aligned} \text{DFT}(\bar{\mathbf{z}}) &= \begin{bmatrix} \text{DFT}(\bar{\mathbf{z}})[1]^* \\ \vdots \\ \text{DFT}(\bar{\mathbf{z}})[C]^* \end{bmatrix} = [\text{DFT}(\bar{\mathbf{z}})(0), \text{DFT}(\bar{\mathbf{z}})(1), \dots, \text{DFT}(\bar{\mathbf{z}})(T-1)] \\ &= \{\text{DFT}(\bar{\mathbf{z}})[c](t)\}_{c=1, t=0}^{c=C, t=T-1}. \end{aligned} \quad (45)$$

As such, we have $\text{DFT}(\mathbf{z}[c]) = \text{DFT}(\bar{\mathbf{z}})[c]$.

By using Fact 6, $\text{circ}(\bar{\mathbf{z}})$ and $\text{DFT}(\bar{\mathbf{z}})$ are related as follows:

$$\text{circ}(\bar{\mathbf{z}}) = \begin{bmatrix} \mathbf{F}_T^* \cdot \text{diag}(\text{DFT}(\mathbf{z}[1])) \cdot \mathbf{F}_T \\ \vdots \\ \mathbf{F}_T^* \cdot \text{diag}(\text{DFT}(\mathbf{z}[C])) \cdot \mathbf{F}_T \end{bmatrix} = \begin{bmatrix} \mathbf{F}_T^* & \cdots & \mathbf{0} \\ \vdots & \ddots & \vdots \\ \mathbf{0} & \cdots & \mathbf{F}_T^* \end{bmatrix} \cdot \begin{bmatrix} \text{diag}(\text{DFT}(\mathbf{z}[1])) \\ \vdots \\ \text{diag}(\text{DFT}(\mathbf{z}[C])) \end{bmatrix} \cdot \mathbf{F}_T. \quad (46)$$

We now explain how this relationship can be leveraged to produce a fast computation of $\bar{\mathbf{E}}$ defined in (31). First, there exists a permutation matrix \mathbf{P} such that

$$\begin{bmatrix} \text{diag}(\text{DFT}(\mathbf{z}[1])) \\ \text{diag}(\text{DFT}(\mathbf{z}[2])) \\ \vdots \\ \text{diag}(\text{DFT}(\mathbf{z}[C])) \end{bmatrix} = \mathbf{P} \cdot \begin{bmatrix} \text{DFT}(\bar{\mathbf{z}})(0) & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \text{DFT}(\bar{\mathbf{z}})(1) & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & \text{DFT}(\bar{\mathbf{z}})(T-1) \end{bmatrix}. \quad (47)$$

Combining (46) and (47), we have

$$\text{circ}(\bar{\mathbf{z}}) \cdot \text{circ}(\bar{\mathbf{z}})^* = \begin{bmatrix} \mathbf{F}_T^* & \cdots & \mathbf{0} \\ \vdots & \ddots & \vdots \\ \mathbf{0} & \cdots & \mathbf{F}_T^* \end{bmatrix} \cdot \mathbf{P} \cdot \mathbf{D}(\bar{\mathbf{z}}) \cdot \mathbf{P}^* \cdot \begin{bmatrix} \mathbf{F}_T & \cdots & \mathbf{0} \\ \vdots & \ddots & \vdots \\ \mathbf{0} & \cdots & \mathbf{F}_T \end{bmatrix}, \quad (48)$$

where

$$\mathbf{D}(\bar{\mathbf{z}}) \doteq \begin{bmatrix} \text{DFT}(\bar{\mathbf{z}})(0) \cdot \text{DFT}(\bar{\mathbf{z}})(0)^* & \cdots & \mathbf{0} \\ \vdots & \ddots & \vdots \\ \mathbf{0} & \cdots & \text{DFT}(\bar{\mathbf{z}})(T-1) \cdot \text{DFT}(\bar{\mathbf{z}})(T-1)^* \end{bmatrix}. \quad (49)$$

Now, consider a collection of m multi-channel 1D signals $\bar{\mathbf{Z}} \in \mathbb{R}^{C \times T \times m}$. Similar to the notation in (28), we denote

$$\begin{aligned} \text{DFT}(\bar{\mathbf{Z}})[c] &= [\text{DFT}(\bar{\mathbf{z}}^1)[c], \dots, \text{DFT}(\bar{\mathbf{z}}^m)[c]] \in \mathbb{C}^{T \times m}, \\ \text{DFT}(\bar{\mathbf{Z}})(p) &= [\text{DFT}(\bar{\mathbf{z}}^1)(p), \dots, \text{DFT}(\bar{\mathbf{z}}^m)(p)] \in \mathbb{C}^{C \times m}. \end{aligned} \quad (50)$$

By using (48), we have

$$\bar{\mathbf{E}} = \begin{bmatrix} \mathbf{F}_T^* & \cdots & \mathbf{0} \\ \vdots & \ddots & \vdots \\ \mathbf{0} & \cdots & \mathbf{F}_T^* \end{bmatrix} \cdot \mathbf{P} \cdot \alpha \cdot \left[\mathbf{I} + \alpha \cdot \sum_{i=1}^m \mathbf{D}(\bar{\mathbf{z}}^i) \right]^{-1} \cdot \mathbf{P}^* \cdot \begin{bmatrix} \mathbf{F}_T & \cdots & \mathbf{0} \\ \vdots & \ddots & \vdots \\ \mathbf{0} & \cdots & \mathbf{F}_T \end{bmatrix}. \quad (51)$$

Note that $\alpha \cdot [\mathbf{I} + \alpha \cdot \sum_{i=1}^m \mathbf{D}(\bar{\mathbf{z}}^i)]^{-1}$ is equal to

$$\begin{aligned} & \alpha \begin{bmatrix} \mathbf{I} + \alpha \text{DFT}(\bar{\mathbf{Z}})(0) \cdot \text{DFT}(\bar{\mathbf{Z}})(0)^* & \cdots & \mathbf{0} \\ \vdots & \ddots & \vdots \\ \mathbf{0} & \cdots & \mathbf{I} + \alpha \text{DFT}(\bar{\mathbf{Z}})(T-1) \cdot \text{DFT}(\bar{\mathbf{Z}})(T-1)^* \end{bmatrix}^{-1} \\ &= \begin{bmatrix} \alpha (\mathbf{I} + \alpha \text{DFT}(\bar{\mathbf{Z}})(0) \cdot \text{DFT}(\bar{\mathbf{Z}})(0)^*)^{-1} & \cdots & \mathbf{0} \\ \vdots & \ddots & \vdots \\ \mathbf{0} & \cdots & \alpha (\mathbf{I} + \alpha \text{DFT}(\bar{\mathbf{Z}})(T-1) \cdot \text{DFT}(\bar{\mathbf{Z}})(T-1)^*)^{-1} \end{bmatrix}. \end{aligned} \quad (52)$$

Therefore, the calculation of $\bar{\mathbf{E}}$ only requires inverting T matrices of size $C \times C$. This motivates us to construct the ReduNet in the spectral domain for the purpose of accelerating the computation, as we explain next.

Shift-invariant ReduNet in the Spectral Domain. Motivated by the result in (52), we introduce the notations $\bar{\mathcal{E}} \in \mathbb{R}^{C \times C \times T}$ and $\bar{\mathcal{C}}^j \in \mathbb{R}^{C \times C \times T}$ given by

$$\bar{\mathcal{E}}(p) \doteq \alpha \cdot [\mathbf{I} + \alpha \cdot \text{DFT}(\bar{\mathbf{Z}})(p) \cdot \text{DFT}(\bar{\mathbf{Z}})(p)^*]^{-1} \in \mathbb{C}^{C \times C}, \quad (53)$$

$$\bar{\mathcal{C}}^j(p) \doteq \alpha_j \cdot [\mathbf{I} + \alpha_j \cdot \text{DFT}(\bar{\mathbf{Z}})(p) \cdot \mathbf{\Pi}_j \cdot \text{DFT}(\bar{\mathbf{Z}})(p)^*]^{-1} \in \mathbb{C}^{C \times C}. \quad (54)$$

In above, $\bar{\mathcal{E}}(p)$ (resp., $\bar{\mathcal{C}}^j(p)$) is the p -th slice of $\bar{\mathcal{E}}$ (resp., $\bar{\mathcal{C}}^j$) on the last dimension. Then, the gradient of $\Delta R_{\text{circ}}(\bar{\mathbf{Z}}, \mathbf{\Pi})$ with respect to $\bar{\mathbf{Z}}$ can be calculated by the following result.

Theorem B.3 (Computing multi-channel convolutions \bar{E} and \bar{C}^j) Let $\bar{U} \in \mathbb{C}^{C \times T \times m}$ and $\bar{W}^j \in \mathbb{C}^{C \times T \times m}, j = 1, \dots, k$ be given by

$$\bar{U}(p) \doteq \bar{\mathcal{E}}(p) \cdot \text{DFT}(\bar{\mathbf{Z}})(p), \quad (55)$$

$$\bar{W}^j(p) \doteq \bar{\mathcal{C}}^j(p) \cdot \text{DFT}(\bar{\mathbf{Z}})(p), \quad j = 1, \dots, k, \quad (56)$$

for each $p \in \{0, \dots, T-1\}$. Then, we have

$$\frac{1}{2T} \frac{\partial \log \det(\mathbf{I} + \alpha \cdot \text{circ}(\bar{\mathbf{Z}}) \text{circ}(\bar{\mathbf{Z}})^*)}{\partial \bar{\mathbf{Z}}} = \text{IDFT}(\bar{\mathbf{U}}), \quad (57)$$

$$\frac{\gamma_j}{2T} \frac{\partial \log \det(\mathbf{I} + \alpha_j \cdot \text{circ}(\bar{\mathbf{Z}}) \bar{\mathbf{\Pi}}^j \text{circ}(\bar{\mathbf{Z}})^*)}{\partial \bar{\mathbf{Z}}} = \gamma_j \cdot \text{IDFT}(\bar{\mathbf{W}}^j \bar{\mathbf{\Pi}}^j). \quad (58)$$

By this result, the gradient ascent update in (3) (when applied to $\Delta R_{\text{circ}}(\bar{\mathbf{Z}}, \bar{\mathbf{\Pi}})$) can be equivalently expressed as an update in frequency domain on $\bar{\mathbf{V}}_\ell \doteq \text{DFT}(\bar{\mathbf{Z}}_\ell)$ as

$$\bar{\mathbf{V}}_{\ell+1}(p) \propto \bar{\mathbf{V}}_\ell(p) + \eta \left(\bar{\mathcal{E}}_\ell(p) \cdot \bar{\mathbf{V}}_\ell(p) - \sum_{j=1}^k \gamma_j \bar{\mathcal{C}}_\ell^j(p) \cdot \bar{\mathbf{V}}_\ell(p) \bar{\mathbf{\Pi}}^j \right), \quad p = 0, \dots, T-1. \quad (59)$$

Similarly, the gradient-guided feature map increment in (13) can be equivalently expressed as an update in frequency domain on $\bar{\mathbf{v}}_\ell \doteq \text{DFT}(\bar{\mathbf{z}}_\ell)$ as

$$\bar{\mathbf{v}}_{\ell+1}(p) \propto \bar{\mathbf{v}}_\ell(p) + \eta \cdot \bar{\mathcal{E}}_\ell(p) \bar{\mathbf{v}}_\ell(p) - \eta \cdot \boldsymbol{\sigma} \left([\bar{\mathcal{C}}_\ell^1(p) \bar{\mathbf{v}}_\ell(p), \dots, \bar{\mathcal{C}}_\ell^k(p) \bar{\mathbf{v}}_\ell(p)] \right), \quad p = 0, \dots, T-1, \quad (60)$$

subject to the constraint that $\|\bar{\mathbf{v}}_{\ell+1}\|_F = \|\bar{\mathbf{z}}_{\ell+1}\|_F = 1$ (the first equality follows from Fact 7).

We summarize the training, or construction to be more precise, of ReduNet in the spectral domain in Algorithm 1.

Proof: [Proof to Theorem (B.3)]

From (4), (51) and (46), we have

$$\frac{1}{2} \frac{\partial \log \det(\mathbf{I} + \alpha \text{circ}(\bar{\mathbf{Z}}) \text{circ}(\bar{\mathbf{Z}})^*)}{\partial \text{circ}(\bar{\mathbf{z}}^i)} = \bar{\mathbf{E}} \text{circ}(\bar{\mathbf{z}}^i) = \bar{\mathbf{E}} \begin{bmatrix} \mathbf{F}_T^* & \cdots & \mathbf{0} \\ \vdots & \ddots & \vdots \\ \mathbf{0} & \cdots & \mathbf{F}_T^* \end{bmatrix} \begin{bmatrix} \text{diag}(\text{DFT}(\mathbf{z}^i[1])) \\ \vdots \\ \text{diag}(\text{DFT}(\mathbf{z}^i[C])) \end{bmatrix} \mathbf{F}_T \quad (61)$$

$$= \begin{bmatrix} \mathbf{F}_T^* & \cdots & \mathbf{0} \\ \vdots & \ddots & \vdots \\ \mathbf{0} & \cdots & \mathbf{F}_T^* \end{bmatrix} \cdot \mathbf{P} \cdot \alpha \cdot \left[\mathbf{I} + \alpha \cdot \sum_i \mathbf{D}(\bar{\mathbf{z}}^i) \right]^{-1} \cdot \begin{bmatrix} \text{DFT}(\bar{\mathbf{z}}^i)(0) & \cdots & \mathbf{0} \\ \vdots & \ddots & \vdots \\ \mathbf{0} & \cdots & \text{DFT}(\bar{\mathbf{z}}^i)(T-1) \end{bmatrix} \cdot \mathbf{F}_T \quad (62)$$

$$= \begin{bmatrix} \mathbf{F}_T^* & \cdots & \mathbf{0} \\ \vdots & \ddots & \vdots \\ \mathbf{0} & \cdots & \mathbf{F}_T^* \end{bmatrix} \cdot \mathbf{P} \cdot \begin{bmatrix} \bar{\mathcal{E}}(0) \cdot \text{DFT}(\bar{\mathbf{z}}^i)(0) & \cdots & \mathbf{0} \\ \vdots & \ddots & \vdots \\ \mathbf{0} & \cdots & \bar{\mathcal{E}}(T-1) \cdot \text{DFT}(\bar{\mathbf{z}}^i)(T-1) \end{bmatrix} \cdot \mathbf{F}_T \quad (63)$$

$$= \begin{bmatrix} \mathbf{F}_T^* & \cdots & \mathbf{0} \\ \vdots & \ddots & \vdots \\ \mathbf{0} & \cdots & \mathbf{F}_T^* \end{bmatrix} \cdot \mathbf{P} \cdot \begin{bmatrix} \bar{\mathbf{u}}^i(0) & \cdots & \mathbf{0} \\ \vdots & \ddots & \vdots \\ \mathbf{0} & \cdots & \bar{\mathbf{u}}^i(T-1) \end{bmatrix} \cdot \mathbf{F}_T = \begin{bmatrix} \mathbf{F}_T^* & \cdots & \mathbf{0} \\ \vdots & \ddots & \vdots \\ \mathbf{0} & \cdots & \mathbf{F}_T^* \end{bmatrix} \cdot \begin{bmatrix} \text{diag}(\bar{\mathbf{u}}^i[1]) \\ \vdots \\ \text{diag}(\bar{\mathbf{u}}^i[C]) \end{bmatrix} \cdot \mathbf{F}_T \quad (64)$$

$$= \text{circ}(\text{IDFT}(\bar{\mathbf{u}}^i)). \quad (65)$$

Therefore, we have

$$\frac{1}{2} \frac{\partial \log \det(\mathbf{I} + \alpha \cdot \text{circ}(\bar{\mathbf{Z}}) \cdot \text{circ}(\bar{\mathbf{Z}})^*)}{\partial \bar{\mathbf{z}}^i} = \frac{1}{2} \frac{\partial \log \det(\mathbf{I} + \alpha \cdot \text{circ}(\bar{\mathbf{Z}}) \cdot \text{circ}(\bar{\mathbf{Z}})^*)}{\partial \text{circ}(\bar{\mathbf{z}}^i)} \cdot \frac{\partial \text{circ}(\bar{\mathbf{z}}^i)}{\partial \bar{\mathbf{z}}^i} = T \cdot \text{IDFT}(\bar{\mathbf{u}}^i). \quad (66)$$

By collecting the results for all i , we have

$$\frac{\partial \frac{1}{2T} \log \det(\mathbf{I} + \alpha \cdot \text{circ}(\bar{\mathbf{Z}}) \cdot \text{circ}(\bar{\mathbf{Z}})^*)}{\partial \bar{\mathbf{Z}}} = \text{IDFT}(\bar{\mathbf{U}}). \quad (67)$$

Algorithm 1 Training Algorithm (1D Signal, Shift Invariance, Spectral Domain)

Input: $\bar{\mathbf{Z}} \in \mathbb{R}^{C \times T \times m}$, $\mathbf{\Pi}$, $\epsilon > 0$, λ , and a learning rate η .

```

1: Set  $\alpha = \frac{C}{m\epsilon^2}$ ,  $\{\alpha_j = \frac{C}{\text{tr}(\mathbf{\Pi}^j)\epsilon^2}\}_{j=1}^k$ ,  $\{\gamma_j = \frac{\text{tr}(\mathbf{\Pi}^j)}{m}\}_{j=1}^k$ .
2: Set  $\bar{\mathbf{V}}_0 = \{\bar{\mathbf{v}}_0^i(p) \in \mathbb{C}^C\}_{p=0, i=1}^{T-1, m} \doteq \text{DFT}(\bar{\mathbf{Z}}) \in \mathbb{C}^{C \times T \times m}$ .
3: for  $\ell = 1, 2, \dots, L$  do
4:   # Step 1: Compute  $\mathcal{E}$  and  $\mathcal{C}$ .
5:   for  $p = 0, 1, \dots, T-1$  do
6:     Compute  $\bar{\mathcal{E}}_\ell(p) \in \mathbb{C}^{C \times C}$  and  $\{\bar{\mathcal{C}}_\ell^j(p) \in \mathbb{C}^{C \times C}\}_{j=1}^k$  as
       
$$\bar{\mathcal{E}}_\ell(p) \doteq \alpha \cdot [\mathbf{I} + \alpha \cdot \bar{\mathbf{V}}_{\ell-1}(p) \cdot \bar{\mathbf{V}}_{\ell-1}(p)^*]^{-1},$$

       
$$\bar{\mathcal{C}}_\ell^j(p) \doteq \alpha_j \cdot [\mathbf{I} + \alpha_j \cdot \bar{\mathbf{V}}_{\ell-1}(p) \cdot \mathbf{\Pi}^j \cdot \bar{\mathbf{V}}_{\ell-1}(p)^*]^{-1};$$

7:   end for
8:   # Step 2: Update  $\bar{\mathbf{v}}^i$  for each  $i$ .
9:   for  $i = 1, \dots, m$  do
10:    # Compute projection at each frequency  $p$ .
11:    for  $p = 0, 1, \dots, T-1$  do
12:      Compute  $\{\bar{\mathbf{p}}_\ell^{ij}(p) \doteq \bar{\mathcal{C}}_\ell^j(p) \cdot \bar{\mathbf{v}}_\ell^i(p) \in \mathbb{C}^{C \times 1}\}_{j=1}^k$ ;
13:    end for
14:    # Compute overall projection by aggregating over frequency  $p$ .
15:    Let  $\{\bar{\mathbf{P}}_\ell^{ij} = [\bar{\mathbf{p}}_\ell^{ij}(0), \dots, \bar{\mathbf{p}}_\ell^{ij}(T-1)] \in \mathbb{C}^{C \times T}\}_{j=1}^k$ ;
16:    # Compute soft assignment from projection.
17:    Compute  $\{\hat{\pi}_\ell^{ij} = \frac{\exp(-\lambda \|\bar{\mathbf{P}}_\ell^{ij}\|_F)}{\sum_{j=1}^k \exp(-\lambda \|\bar{\mathbf{P}}_\ell^{ij}\|_F)}\}_{j=1}^k$ ;
18:    # Compute update at each frequency  $p$ .
19:    for  $p = 0, 1, \dots, T-1$  do
20:      
$$\bar{\mathbf{v}}_\ell^i(p) = \bar{\mathbf{v}}_{\ell-1}^i(p) + \eta \left( \bar{\mathcal{E}}_\ell(p) \bar{\mathbf{v}}_\ell^i(p) - \sum_{j=1}^k \gamma_j \cdot \hat{\pi}_\ell^{ij} \cdot \bar{\mathcal{C}}_\ell^j(p) \cdot \bar{\mathbf{v}}_\ell^i(p) \right);$$

21:    end for
22:    Normalize  $\bar{\mathbf{v}}_\ell^i = \bar{\mathbf{v}}_\ell^i / \|\bar{\mathbf{v}}_\ell^i\|_F$ ;
23:  end for
24:  Set  $\bar{\mathbf{Z}}_\ell = \text{IDFT}(\bar{\mathbf{V}}_\ell)$  as the feature at the  $\ell$ -th layer;
25:  # Evaluate the objective value.
26:  
$$\frac{1}{2T} \sum_{p=0}^{T-1} \left( \log \det[\mathbf{I} + \alpha \bar{\mathbf{V}}_\ell(p) \cdot \bar{\mathbf{V}}_\ell(p)^*] - \frac{\text{tr}(\mathbf{\Pi}^j)}{m} \log \det[\mathbf{I} + \alpha_j \bar{\mathbf{V}}_\ell(p) \cdot \mathbf{\Pi}^j \cdot \bar{\mathbf{V}}_\ell(p)^*] \right);$$

27: end for
Output: features  $\bar{\mathbf{Z}}_L$ , the learned filters  $\{\bar{\mathcal{E}}_\ell(p)\}_{\ell, p}$  and  $\{\bar{\mathcal{C}}_\ell^j(p)\}_{j, \ell, p}$ .
```

In a similar fashion, we get

$$\frac{\partial \frac{\gamma_j}{2T} \log \det \left(\mathbf{I} + \alpha_j \cdot \text{circ}(\bar{\mathbf{Z}}) \cdot \mathbf{\Pi}^j \cdot \text{circ}(\bar{\mathbf{Z}})^* \right)}{\partial \bar{\mathbf{Z}}} = \gamma_j \cdot \text{IDFT}(\bar{\mathbf{W}}^j \cdot \mathbf{\Pi}^j). \quad (68)$$

□

C 2D CIRCULAR TRANSLATION INVARIANCE

To a large degree, both conceptually and technically, the 2D case is very similar to the 1D case that we have studied carefully in the previous Appendix B. For the sake of consistency and completeness, we here give a brief account.

C.1 DOUBLY BLOCK CIRCULANT MATRIX

In this section, we consider z as a 2D signal such as an image, and use H and W to denote its “height” and “width”, respectively. It will be convenient to work with both a matrix representation

$$z = \begin{bmatrix} z(0,0) & z(0,1) & \cdots & z(0,W-1) \\ z(1,0) & z(1,1) & \cdots & z(1,W-1) \\ \vdots & \vdots & \ddots & \vdots \\ z(H-1,0) & z(H-1,1) & \cdots & z(H-1,W-1) \end{bmatrix} \in \mathbb{R}^{H \times W}, \quad (69)$$

as well as a vector representation

$$\text{vec}(z) \doteq \begin{bmatrix} z(0,0), \dots, z(0,W-1), z(1,0), \dots, z(1,W-1), \dots \\ \dots, z(H-1,0), \dots, z(H-1,W-1) \end{bmatrix}^* \in \mathbb{R}^{(H \times W)}. \quad (70)$$

We represent the circular translated version of z as $\text{trans}_{p,q}(z) \in \mathbb{R}^{H \times W}$ by an amount of p and q on the vertical and horizontal directions, respectively. That is, we let

$$\text{trans}_{p,q}(z)(h,w) \doteq z(h-p \bmod H, w-q \bmod W), \quad \forall (h,w) \in \{0, \dots, H-1\} \times \{0, \dots, W-1\}. \quad (71)$$

It is obvious that $\text{trans}_{0,0}(z) = z$. Moreover, there is a total number of $H \times W$ distinct translations given by $\{\text{trans}_{p,q}(z), (p,q) \in \{0, \dots, H-1\} \times \{0, \dots, W-1\}\}$. We may arrange the vector representations of them into a matrix and obtain

$$\begin{aligned} \text{circ}(z) \doteq & \begin{bmatrix} \text{vec}(\text{trans}_{0,0}(z)), \dots, \text{vec}(\text{trans}_{0,W-1}(z)), \\ \text{vec}(\text{trans}_{1,0}(z)), \dots, \text{vec}(\text{trans}_{1,W-1}(z)), \\ \dots, \\ \text{vec}(\text{trans}_{H-1,0}(z)), \dots, \text{vec}(\text{trans}_{H-1,W-1}(z)) \end{bmatrix} \in \mathbb{R}^{(H \times W) \times (H \times W)}. \end{aligned} \quad (72)$$

The matrix $\text{circ}(z)$ is known as the *doubly block circulant matrix* associated with z (see, e.g., [Abidi et al. \(2016\)](#); [Sedghi et al. \(2018\)](#)).

We now consider a multi-channel 2D signal represented as a tensor $\bar{z} \in \mathbb{R}^{C \times H \times W}$, where C is the number of channels. The c -th channel of \bar{z} is represented as $\bar{z}[c] \in \mathbb{R}^{H \times W}$, and the (h,w) -th pixel is represented as $\bar{z}(h,w) \in \mathbb{R}^C$. To compute the coding rate reduction for a collection of such multi-channel 2D signals, we may flatten the tensor representation into a vector representation by concatenating the vector representation of each channel, i.e., we let

$$\text{vec}(\bar{z}) = [\text{vec}(\bar{z}[1])^*, \dots, \text{vec}(\bar{z}[C])^*]^* \in \mathbb{R}^{(C \times H \times W)} \quad (73)$$

Furthermore, to obtain shift invariance for coding rate reduction, we may generate a collection of translated versions of \bar{z} (along two spatial dimensions). Stacking the vector representation for such translated copies as column vectors, we obtain

$$\text{circ}(\bar{z}) \doteq \begin{bmatrix} \text{circ}(\bar{z}[1]) \\ \vdots \\ \text{circ}(\bar{z}[C]) \end{bmatrix} \in \mathbb{R}^{(C \times H \times W) \times (H \times W)}. \quad (74)$$

We can now define a *translation invariant coding rate reduction* for multi-channel 2D signals. Consider a collection of m multi-channel 2D signals $\{\bar{z}^i \in \mathbb{R}^{C \times H \times W}\}_{i=1}^m$. Compactly representing the data by $\bar{Z} \in \mathbb{R}^{C \times H \times W \times m}$ where the i -th slice on the last dimension is \bar{z}^i , we denote

$$\text{circ}(\bar{Z}) = [\text{circ}(\bar{z}^1), \dots, \text{circ}(\bar{z}^m)] \in \mathbb{R}^{(C \times H \times W) \times (H \times W \times m)}. \quad (75)$$

Then, we define

$$\Delta R_{\text{circ}}(\bar{\mathbf{Z}}, \mathbf{\Pi}) \doteq \frac{1}{HW} \Delta R(\text{circ}(\bar{\mathbf{Z}}), \bar{\mathbf{\Pi}}) = \frac{1}{2HW} \log \det \left(\mathbf{I} + \alpha \cdot \text{circ}(\bar{\mathbf{Z}}) \cdot \text{circ}(\bar{\mathbf{Z}})^* \right) - \sum_{j=1}^k \frac{\gamma_j}{2HW} \log \det \left(\mathbf{I} + \alpha_j \cdot \text{circ}(\bar{\mathbf{Z}}) \cdot \bar{\mathbf{\Pi}}^j \cdot \text{circ}(\bar{\mathbf{Z}})^* \right), \quad (76)$$

where $\alpha = \frac{CHW}{mHW\epsilon^2} = \frac{C}{m\epsilon^2}$, $\alpha_j = \frac{CHW}{\text{tr}(\mathbf{\Pi}^j)HW\epsilon^2} = \frac{C}{\text{tr}(\mathbf{\Pi}^j)\epsilon^2}$, $\gamma_j = \frac{\text{tr}(\mathbf{\Pi}^j)}{m}$, and $\bar{\mathbf{\Pi}}^j$ is augmented membership matrix in an obvious way.

By following an analogous argument as in the 1D case, one can show that ReduNet for multi-channel 2D signals naturally gives rise to the multi-channel 2D circulant convolution operations. We omit the details, and focus on the construction of ReduNet in the frequency domain.

C.2 FAST COMPUTATION IN SPECTRAL DOMAIN

Doubly block circulant matrix and 2D-DFT. Similar to the case of circulant matrices for 1D signals, all doubly block circulant matrices share the same set of eigenvectors, and these eigenvectors form a unitary matrix given by

$$\mathbf{F} \doteq \mathbf{F}_H \otimes \mathbf{F}_W \in \mathbb{C}^{(H \times W) \times (H \times W)}, \quad (77)$$

where \otimes denotes the Kronecker product and $\mathbf{F}_H, \mathbf{F}_W$ are defined as in (36).

Analogous to Fact 4, \mathbf{F} defines 2D-DFT as follows.

Fact 8 (2D-DFT as matrix-vector multiplication) *The 2D-DFT of a signal $\mathbf{z} \in \mathbb{R}^{H \times W}$ can be computed as*

$$\text{vec}(\text{DFT}(\mathbf{z})) \doteq \mathbf{F} \cdot \text{vec}(\mathbf{z}) \in \mathbb{C}^{(H \times W)}, \quad (78)$$

where

$$\text{DFT}(\mathbf{z})(p, q) = \frac{1}{\sqrt{H \cdot W}} \sum_{h=0}^{H-1} \sum_{w=0}^{W-1} \mathbf{z}(h, w) \cdot \omega_H^{p \cdot h} \omega_W^{q \cdot w}, \quad \forall (p, q) \in \{0, \dots, H-1\} \times \{0, \dots, W-1\}. \quad (79)$$

The 2D-IDFT of a signal $\mathbf{v} \in \mathbb{C}^{H \times W}$ can be computed as

$$\text{vec}(\text{IDFT}(\mathbf{v})) \doteq \mathbf{F}_T^* \cdot \text{vec}(\mathbf{v}) \in \mathbb{C}^{(H \times W)}, \quad (80)$$

where

$$\text{IDFT}(\mathbf{v})(h, w) = \frac{1}{\sqrt{H \cdot W}} \sum_{p=0}^{H-1} \sum_{q=0}^{W-1} v(p, q) \cdot \omega_H^{-p \cdot h} \omega_W^{-q \cdot w}, \quad \forall (h, w) \in \{0, \dots, H-1\} \times \{0, \dots, W-1\}. \quad (81)$$

Analogous to Fact 9, \mathbf{F} relates $\text{DFT}(\mathbf{z})$ and $\text{circ}(\mathbf{z})$ as follows.

Fact 9 (2D-DFT are eigenvalues of the doubly block circulant matrix) *Given a signal $\mathbf{z} \in \mathbb{C}^{H \times W}$, we have*

$$\mathbf{F} \cdot \text{circ}(\mathbf{z}) \cdot \mathbf{F}^* = \text{diag}(\text{vec}(\text{DFT}(\mathbf{z}))) \quad \text{or} \quad \text{circ}(\mathbf{z}) = \mathbf{F}^* \cdot \text{diag}(\text{vec}(\text{DFT}(\mathbf{z}))) \cdot \mathbf{F}. \quad (82)$$

Doubly block circulant matrix and 2D-DFT for multi-channel signals. We now consider multi-channel 2D signals $\bar{z} \in \mathbb{R}^{C \times H \times W}$. Let $\text{DFT}(\bar{z}) \in \mathbb{C}^{C \times H \times W}$ be a matrix where the c -th slice on the first dimension is the DFT of the corresponding signal $z[c]$. That is, $\text{DFT}(\bar{z})[c] = \text{DFT}(z[c]) \in \mathbb{C}^{H \times W}$. We use $\text{DFT}(\bar{z})(p, q) \in \mathbb{C}^C$ to denote slicing of \bar{z} on the frequency dimensions.

By using Fact 9, $\text{circ}(\bar{z})$ and $\text{DFT}(\bar{z})$ are related as follows:

$$\begin{aligned} \text{circ}(\bar{z}) &= \begin{bmatrix} \mathbf{F}^* \cdot \text{diag}(\text{vec}(\text{DFT}(z[1]))) \cdot \mathbf{F} \\ \vdots \\ \mathbf{F}^* \cdot \text{diag}(\text{vec}(\text{DFT}(z[C]))) \cdot \mathbf{F} \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{F}^* & \cdots & \mathbf{0} \\ \mathbf{0} & \cdots & \mathbf{0} \\ \vdots & \ddots & \vdots \\ \mathbf{0} & \cdots & \mathbf{F}^* \end{bmatrix} \cdot \begin{bmatrix} \text{diag}(\text{vec}(\text{DFT}(z[1]))) \\ \text{diag}(\text{vec}(\text{DFT}(z[2]))) \\ \vdots \\ \text{diag}(\text{vec}(\text{DFT}(z[C]))) \end{bmatrix} \cdot \mathbf{F}. \end{aligned} \quad (83)$$

Similar to the 1D case, this relation can be leveraged to produce a fast implementation of ReduNet in the spectral domain.

Translation-invariant ReduNet in the Spectral Domain. Given a collection of multi-channel 2D signals $\bar{\mathbf{Z}} \in \mathbb{R}^{C \times H \times W \times m}$, we denote

$$\text{DFT}(\bar{\mathbf{Z}})(p, q) \doteq [\text{DFT}(\bar{z}^1)(p, q), \dots, \text{DFT}(\bar{z}^m)(p, q)] \in \mathbb{R}^{C \times m}. \quad (84)$$

We introduce the notations $\bar{\mathcal{E}}(p, q) \in \mathbb{R}^{C \times C \times H \times W}$ and $\bar{\mathcal{C}}^j(p, q) \in \mathbb{R}^{C \times C \times H \times W}$ given by

$$\bar{\mathcal{E}}(p, q) \doteq \alpha \cdot [\mathbf{I} + \alpha \cdot \text{DFT}(\bar{\mathbf{Z}})(p, q) \cdot \text{DFT}(\bar{\mathbf{Z}})(p, q)^*]^{-1} \in \mathbb{C}^{C \times C}, \quad (85)$$

$$\bar{\mathcal{C}}^j(p, q) \doteq \alpha_j \cdot [\mathbf{I} + \alpha_j \cdot \text{DFT}(\bar{\mathbf{Z}})(p, q) \cdot \mathbf{\Pi}_j \cdot \text{DFT}(\bar{\mathbf{Z}})(p, q)^*]^{-1} \in \mathbb{C}^{C \times C}. \quad (86)$$

In above, $\bar{\mathcal{E}}(p, q)$ (resp., $\bar{\mathcal{C}}^j(p, q)$) is the (p, q) -th slice of $\bar{\mathcal{E}}$ (resp., $\bar{\mathcal{C}}^j$) on the last two dimensions. Then, the gradient of $\Delta R_{\text{circ}}(\bar{\mathbf{Z}}, \mathbf{\Pi})$ with respect to $\bar{\mathbf{Z}}$ can be calculated by the following result.

Theorem C.1 (Computing multi-channel 2D convolutions $\bar{\mathcal{E}}$ and $\bar{\mathcal{C}}^j$) Let $\bar{\mathbf{U}} \in \mathbb{C}^{C \times H \times W \times m}$ and $\bar{\mathbf{W}}^j \in \mathbb{C}^{C \times H \times W \times m}$, $j = 1, \dots, k$ be given by

$$\bar{\mathbf{U}}(p, q) \doteq \bar{\mathcal{E}}(p, q) \cdot \text{DFT}(\bar{\mathbf{Z}})(p, q), \quad (87)$$

$$\bar{\mathbf{W}}^j(p, q) \doteq \bar{\mathcal{C}}^j(p, q) \cdot \text{DFT}(\bar{\mathbf{Z}})(p, q), \quad j = 1, \dots, k, \quad (88)$$

for each $(p, q) \in \{0, \dots, H-1\} \times \{0, \dots, W-1\}$. Then, we have

$$\frac{1}{2HW} \frac{\partial \log \det(\mathbf{I} + \alpha \cdot \text{circ}(\bar{\mathbf{Z}}) \text{circ}(\bar{\mathbf{Z}})^*)}{\partial \bar{\mathbf{Z}}} = \text{IDFT}(\bar{\mathbf{U}}), \quad (89)$$

$$\frac{1}{2HW} \frac{\partial (\gamma_j \log \det(\mathbf{I} + \alpha_j \cdot \text{circ}(\bar{\mathbf{Z}}) \mathbf{\Pi}^j \text{circ}(\bar{\mathbf{Z}})^*))}{\partial \bar{\mathbf{Z}}} = \gamma_j \cdot \text{IDFT}(\bar{\mathbf{W}}^j \mathbf{\Pi}^j). \quad (90)$$

This result shows that the calculation of the derivatives for the 2D case is analogous to that of the 1D case. Therefore, the construction of the ReduNet for 2D translation invariance can be performed using Algorithm 1 with straightforward extensions.

D IMPLEMENTATION DETAILS AND ADDITIONAL EXPERIMENTS

Disclaimer: in this work we do not particularly optimize any of the hyper parameters, such as the number of initial channels, kernel sizes, and learning rate etc., for the best performance. The choices are mostly for convenience and just minimally adequate to verify the concept, due to limited computational resource.

D.1 ADDITIONAL EXPERIMENTS ON LEARNING MIXTURE OF GAUSSIANS IN \mathbb{S}^1 AND \mathbb{S}^2

We provide the cosine similarity results for the experiments described in Figure 4. The results are shown in Figure 8. We can observe that the network can map the data points to orthogonal subspaces.

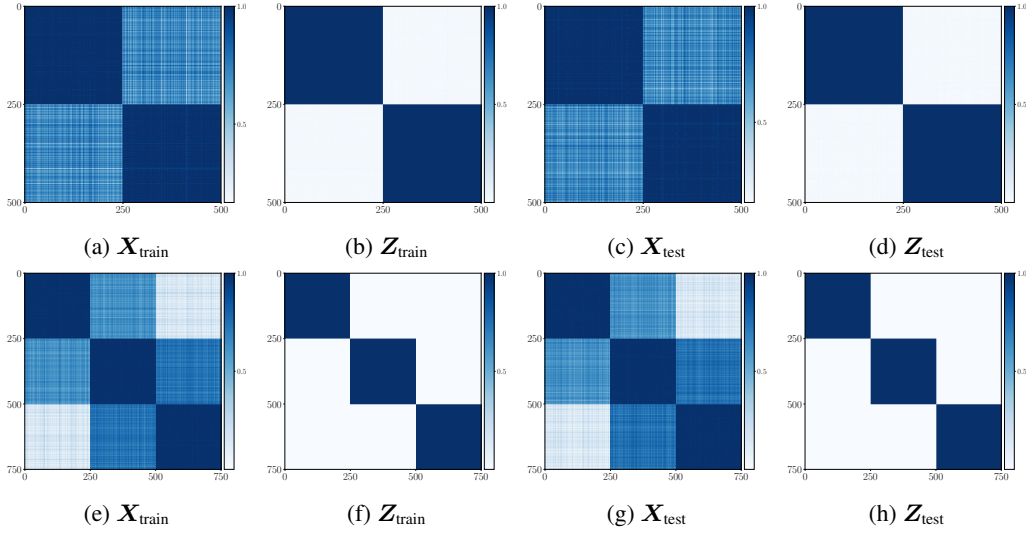


Figure 8: Cosine similarity (absolute value) for 2D and 3D Mixture of Gaussians. Lighter color implies samples are more orthogonal.

Additional experiments on \mathbb{S}^1 and \mathbb{S}^2 . We also provide additional experiments on learning mixture of Gaussians in \mathbb{S}^1 and \mathbb{S}^2 in Figure 9. We can observe similar behavior of the proposed ReduNet: the network can map data points from different classes to orthogonal subspaces.




Figure 9: Learning mixture of Gaussians in \mathbb{S}^1 and \mathbb{S}^2 . For \mathbb{S}^1 , we set $\sigma_1 = \sigma_2 = 0.1$.

Additional experiments on \mathbb{S}^1 with more than 2 classes. We try to apply ReduNet to learn mixture of Gaussian distributions on \mathbb{S}^1 with the number of class is larger than 2. Notice that these are the cases to which the existing theory about MCR² (Yu et al., 2020) no longer applies. These experiments suggest that the MCR² still promotes between-class discriminativeness with so constructed ReduNet. In particular, the case on the left of Figure 11 indicates that the ReduNet has “merged” two linearly correlated clusters into one on the same line. This is consistent with the objective of rate reduction to group data as linear subspaces.




Figure 10: Learning mixture of Gaussians in \mathbb{S}^1 and \mathbb{S}^2 . For \mathbb{S}^2 , we set $\sigma_1 = \sigma_2 = \sigma_3 = 0.1$.




Figure 11: Learning mixture of Gaussian distributions with more than 2 classes. For both cases, we use step size $\eta = 0.5$ and precision $\epsilon = 0.1$. For (a), we set iteration $L = 2,500$; for (b), we set iteration $L = 4,000$.

D.2 EXPERIMENTS ON UCI DATASETS

We evaluate the proposed ReduNet on some real datasets, namely the two UCI tasks (Dua & Graff, 2017): iris and mice. There are 3 classes in iris dataset and the number of features is 4. For mice dataset, there are 8 classes and the number of features is 82. We randomly select 70% data as the training data, and use the rest for evaluation. The results are summarized in Table 2. We compare our method with logistic regression, SVM, and random forest, and we use the implementations by **sklearn** (Pedregosa et al., 2011). From Table 2, we find that the forward-constructed ReduNet is able to achieve comparable performance with classic methods such as logistic regression, SVM, and random forest.

Table 2: Performance (Accuracy) on iris and mice of the UCI datasets.

	REDUNET	LOGISTIC REGRESSION	SVM	RANDOM FOREST
IRIS	0.978	0.933	0.933	0.978
MICE	0.972	0.855	0.975	0.985

D.3 ADDITIONAL EXPERIMENTS ON LEARNING SHIFT INVARIANT FEATURES

We provide additional experiments for *Learning Shift Invariant Features* in §3. The code for sampling from $h_1(t) = \sin(t) + \epsilon$ and $h_2(t) = \text{sign}(\sin(t)) + \epsilon$ is described in Algorithm 2, and the pseudocode for sampling from 2 classes $\{h_1, h_2\}$ is described as follows, we sample training and test signals using the same procedure.

```
t0 = np.random.uniform(low=0, high=10*np.pi, size=samples)
x = np.linspace(t0, t0+2*np.pi, time).T
noise1 = np.random.normal(0, 0.1, size=(samples, time))
X1 = np.sin(x) + noise1
noise2 = np.random.normal(0, 0.1, size=(samples, time))
X2 = np.sign(np.sin(x)) + noise2
data = np.vstack([X1, X2])
labels = np.hstack([np.ones(samples)*1,
                    np.ones(samples)*2]).astype(np.int32)
```

We also provide cosine similarities between samples in Figure 13. We visualize the cosine similarities for the input $\mathbf{X}_{\text{train}}, \mathbf{X}_{\text{test}}$ as well as the learned representations $\mathbf{Z}_{\text{train}}, \mathbf{Z}_{\text{test}}$. The cosine similarity between sample pairs selected from different classes are shown in Figure 14. We can observe that the original data is not orthogonal w.r.t. different classes, and the the ReduNet is able to learn discriminative (orthogonal) representations.

Algorithm 2 Pseudocode for sampling signals from 1D functions

Input: Number of samples m , number of classes k , number of features n , function $\{h_1, \dots, h_k\}$.

```
1: for  $j = 1, 2, \dots, k$  do
2:   for  $i = 1, 2, \dots, m$  do
3:      $t_0 \sim \text{Uniform}[0, 10\pi]$ ;
4:      $\mathbf{t} = [t_0, t_0 + 2\pi/n, t_0 + (2\pi/n) \cdot 2, t_0 + (2\pi/n) \cdot 3, \dots, t_0 + (2\pi/n) \cdot (n-1)]$ ;
5:      $\mathbf{x}_j^i = h_j(\mathbf{t}) + \epsilon$  # broadcast over vector  $\mathbf{t}$ ;
6:   end for
7:    $\mathbf{X}_j = [\mathbf{x}_j^1, \mathbf{x}_j^2, \dots, \mathbf{x}_j^m]$ ;
8: end for
9:  $\mathbf{X} = [\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_k]$ ;
10: shuffle  $\mathbf{X}$ .
```

Output: outputs \mathbf{X} .

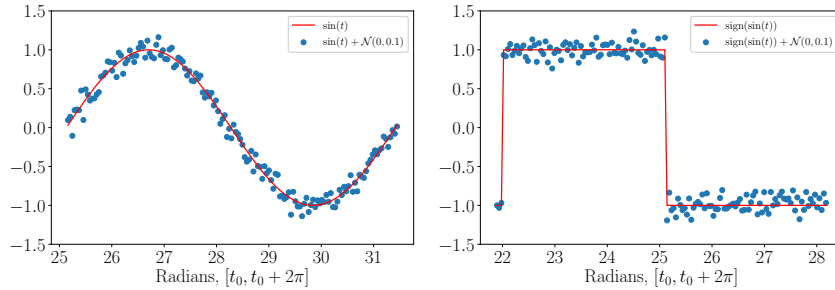


Figure 12: Visualization of signals in 1D. Blue dots represent the sampled signal used for training with dimension $n = 150$. Red curves represent the underlying 1D function (noiseless). **(Left)** One sample from class 1; **(Right)** One sample from class 2.

D.4 ADDITIONAL EXPERIMENTS ON LEARNING ROTATIONAL INVARIANCE ON MNIST

We provide additional experiments for *learning rotational invariance on MNIST* in §3. Examples of rotated images are shown in Figure 17. We also provide cosine similarities between samples in




Figure 13: Cosine similarity (absolute value) of training/test data as well as training/test representations for learning 1D functions.




Figure 14: Histogram of cosine similarity between pairs sampled from different classes for learning 1D function. The histogram of cosine similarity between training data X_c as well as representations Z_c vs. testing (shifted) data $\tilde{X}_{c'}$ as well as (shifted) representations $\tilde{Z}_{c'}$, where we let c denote the class index and $c \neq c'$.

Figure 15. We visualize the cosine similarities for the input $X_{\text{train}}, X_{\text{test}}$ as well as the learned representations $Z_{\text{train}}, Z_{\text{test}}$. The cosine similarity between sample pairs selected from different classes are shown in Figure 16. We can observe that the constructed ReduNet is able to learn discriminative (orthogonal) and invariant representations for MNIST digits.

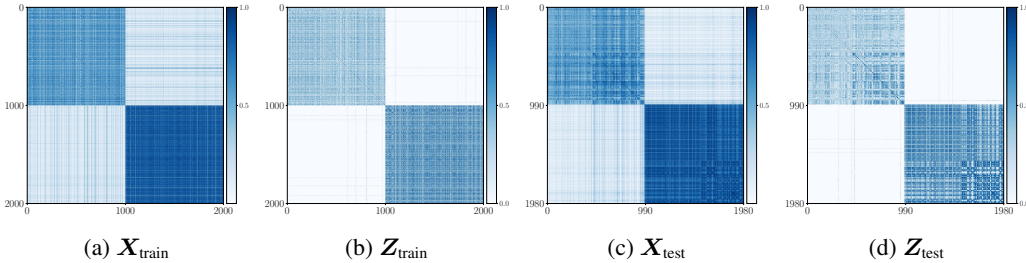


Figure 15: Cosine similarity (absolute value) of training/test data as well as training/test representations for learning rotational invariant representations on MNIST.




Figure 16: Histogram of cosine similarity between pairs sampled from different classes for learning rotational invariant representations on MNIST. The histogram of cosine similarity between training data X_c as well as representations Z_c vs. testing (shifted) data $\bar{X}_{c'}$ as well as (shifted) representations $\bar{Z}_{c'}$, where we let c denote the class index and $c \neq c'$.




Figure 17: Examples of rotated images of MNIST digits for testing rotation invariance, each rotated by 18°. (Left) digit '0'. (Right) digit '1'.